SYSTEMS, METHODS, DEVICES, AND PLATFORMS FOR INDUSTRIAL INTERNET OF THINGS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of priority to the following U.S. Patent Applications: Serial No. 63/625,594, filed January 26, 2024; Serial No. 63/638,596, filed April 25, 2024; and Serial No. 63/639,909, filed April 29, 2024. Each of the above applications is hereby incorporated by reference in its entirety as if fully set forth herein.

BACKGROUND

FIELD

[0002] The present disclosure relates to the field of enterprise management platforms, more particularly involving data management, artificial intelligence, network connectivity and digital twins.

DESCRIPTION OF THE RELATED ART

[0003] Industrial environments, such as environments for large scale manufacturing (such as manufacturing of aircraft, ships, trucks, automobiles, and large industrial machines), energy production environments (such as oil and gas plants, renewable energy environments, and others), energy extraction environments (such as mining, drilling, and the like), construction environments (such as for construction of large buildings), and others, involve highly complex machines, devices and systems and highly complex workflows, in which operators must account for a host of parameters, metrics, and the like in order to optimize design, development, deployment, and operation of different technologies in order to improve overall results. Industrial environments are widely populated with large, complex, heavy machines that are designed to have very relatively long working lifetimes and have ongoing service requirements, including requirements for scheduled maintenance and for often unanticipated repairs. Many of the large industrial machines that require ongoing maintenance, service and repairs are involved in high stakes production processes and other processes, such as energy production, manufacturing, mining, drilling, and transportation, that preferably involve minimal or no interruption. An unanticipated problem, or an extended delay in a service operation that requires a shutdown of a machine that is critical to such a process can cost thousands, or even millions of dollars per day.

[0004] Historically, data has been collected in industrial environments by human beings using dedicated data collectors, often recording batches of specific sensor data on media, such as tape or a hard drive, for later analysis. Batches of data have historically been returned to a central office for analysis, such as undertaking signal processing or other analysis on the data collected by various sensors, after which analysis can be used as a basis for diagnosing problems in an environment and/or suggesting ways to improve operations. This work has historically taken place on a time scale of weeks or months, and has been directed to limited data sets.

[0005] The emergence of the Internet of Things (IoT) has made it possible to connect continuously to, and among, a much wider range of devices. Most such devices are consumer devices, such as lights, thermostats, and the like. With the proliferation of vibration sensors and other Industrial

Internet of Things (IIoT) sensors, there are vast amounts of data available relating to industrial environments. This data is useful in predicting the need for maintenance and for classifying potential issues in the industrial environments. There are, however, many unexplored uses for vibration sensor data and other IIoT sensor data that can improve the operation and uptime of the industrial environments and provide industrial entities with agility in responding to problems before the problems become catastrophic.

[0006] More complex industrial environments remain more difficult, as the range of available data is often limited, and the complexity of dealing with data from multiple sensors makes it much more difficult to produce "smart" solutions that are effective for the industrial sector. For example, in spite of availability of all such data, industrial experts still struggle to properly process all this data because of its sheer size, and thus may not be able to determine faults in the industrial environment when required. A need exists for improved methods and systems for data collection in industrial environments, as well as for improved methods and systems for using collected data to provide improved monitoring, control, intelligent diagnosis of problems and intelligent optimization of operations in various heavy industrial environments.

[0007] Conventional machine vision systems are made of a combination of optics, lighting, sensors and software and aim to replicate the function of the human eye. Such systems typically create an image of an object by capturing and processing the reflected light from the object. An optical lens system typically directs the reflected light to an image sensor device, such as a charge coupled device (CCD) or complementary metal oxide semiconductor (CMOS) device, among others. Such image sensor devices contain arrangements, such as matrices or arrays, of small, accurately spaced photo sensitive elements fabricated using integrated circuit technology. The sensor device converts the light falling on it, through the lens system, into analog electrical signals corresponding to light intensity. The object image is thus broken down into an array of individual picture elements, or pixels. An analog-to-digital converter is used to convert analog voltage outputs of respective. elements into digital values. If the voltage level for each pixel is given either 0 or 1 value depending on whether the analog voltage exceeds some threshold intensity measure, it is called a binary system. In contrast, a gray scale system assigns cardinal values (e.g., in a range of zero to 256), depending on the analog intensity, to each pixel. Thus, in addition to black and white, many different shades of gray can be distinguished. A gray-scale image may be seen to have one channel, represented by a 2-D matrix of pixels having pixel values in the range of, for example, 0 to 255. A color image on the other hand represents the brightness and color of the pixels in an image by the three primary color values: R (red), G (green), and B (blue). Thus, color images have red, green, and blue (RGB) channels, each representing RGB components of the image. This raw data captured by the image sensor is then sent to an image processing system for analysis. The image processing system then processes the raw data to extract useful information to analyze the image and make decisions on such analysis. The image processing system may include a pre-processing function to enhance the image quality. For example, such processing may involve image scaling, noise reduction, color adjustment, brightness adjustment, white balance adjustment, sharpness, adjustment, contrast adjustment and the like. Further the image may be analyzed using machine

learning or other algorithms to identify one or more objects in the image and determine the position and orientation of such objects.

[0008] While vision technology has improved significantly in the past few years, most of the improvements relate to processing of the image data captured by vision sensors and may be attributed to the use of big data, sophisticated machine learning algorithms like convolutional neural networks (CNNs) and graphical processing units (GPUs) for processing of the image data. The conventional vision technology, however, has significant limitations, specifically with respect to capturing of the raw data about an object or a scene. For example, the optical lenses in conventional vision systems attempt to enable extraction of information by relying on focusing techniques that produce images that have good clarity to the human eye. However, the attempt to get an object into focus in fact results in discarding a large amount of relevant information that could otherwise be used in a system, including relevant optical properties.

BRIEF DESCRIPTION OF THE FIGURES

[0009] FIG. 1 through FIG. 5 are diagrammatic views that depict portions of an overall view of an industrial Internet of Things (IoT) data collection, monitoring and control system in accordance with the present disclosure.

[0010] FIG. 6 is a diagrammatic view of a platform including a local data collection system disposed in an industrial environment for collecting data from or about the elements of the environment, such as machines, components, systems, sub-systems, ambient conditions, states, workflows, processes, and other elements in accordance with the present disclosure.

[0011] FIG. 7 is a diagrammatic view that depicts elements of an industrial data collection system for collecting analog sensor data in an industrial environment in accordance with the present disclosure.

[0012] FIG. 8 is a diagrammatic view of a rotating or oscillating machine having a data acquisition module that is configured to collect waveform data in accordance with the present disclosure.

[0013] FIG. 9 is a diagrammatic view of an exemplary tri-axial sensor mounted to a motor bearing of an exemplary rotating machine in accordance with the present disclosure.

[0014] FIG. 10 is a diagrammatic view of components and interactions of a data collection architecture involving application of cognitive and machine learning systems to data collection and processing in accordance with the present disclosure.

[0015] FIG. 11 is a diagrammatic view of components and interactions of a data collection architecture involving application of a platform having a cognitive data marketplace in accordance with the present disclosure.

[0016] FIG. 12 is a diagrammatic view of components and interactions of a data collection architecture involving application of a self-organizing swarm of data collectors in accordance with the present disclosure.

[0017] FIG. 13 is a diagrammatic view of components and interactions of a data collection architecture involving application of a haptic user interface in accordance with the present disclosure.

[0018] FIG. 14 is a diagrammatic view of a multi-format streaming data collection system in accordance with the present disclosure.

[0019] FIG. 15 is a diagrammatic view of combining legacy and streaming data collection and storage in accordance with the present disclosure.

[0020] FIG. 16 is a diagrammatic view of industrial machine sensing using both legacy and updated streamed sensor data processing in accordance with the present disclosure.

[0021] FIG. 17 is a diagrammatic view of an industrial machine sensed data processing system that facilitates portal algorithm use and alignment of legacy and streamed sensor data in accordance with the present disclosure.

[0022] FIG. 18 is a diagrammatic view of components and interactions of a data collection architecture involving a streaming data acquisition instrument receiving analog sensor signals from an industrial environment connected to a cloud network facility in accordance with the present disclosure.

[0023] FIG. 19 is a diagrammatic view of components and interactions of a data collection architecture involving a streaming data acquisition instrument having an alarms module, expert analysis module, and a driver API to facilitate communication with a cloud network facility in accordance with the present disclosure.

[0024] FIG. 20 is a diagrammatic view of components and interactions of a data collection architecture involving a streaming data acquisition instrument and first in, first out memory architecture to provide a real time operating system in accordance with the present disclosure.

[0025] FIG. 21 is a diagrammatic view of components and interactions of a data collection architecture involving a multiple streaming data acquisition instrument receiving analog sensor signals and digitizing those signals to be obtained by a streaming hub server in accordance with the present disclosure.

[0026] FIG. 22 is a diagrammatic view of components and interactions of a data collection architecture involving a master raw data server that processes new streaming data and data already extracted and processed in accordance with the present disclosure.

[0027] FIG. 23, FIG. 24, and FIG. 25 are diagrammatic views of components and interactions of a data collection architecture involving a processing, analysis, report, and archiving server that processes new streaming data and data already extracted and processed in accordance with the present disclosure.

[0028] FIG. 26 is a diagrammatic view of components and interactions of a data collection architecture involving a relation database server and data archives and their connectivity with a cloud network facility in accordance with the present disclosure.

[0029] FIG. 27 through FIG. 32 are diagrammatic views of components and interactions of a data collection architecture involving a virtual streaming data acquisition instrument receiving analog sensor signals from an industrial environment connected to a cloud network facility in accordance with the present disclosure.

[0030] FIG. 33 through FIG. 40 are diagrammatic views of components and interactions of a data collection architecture involving data channel methods and systems for data collection of industrial machines in accordance with the present disclosure.

[0031] FIG. 41 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0032] FIG. 42 and FIG. 43 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0033] FIG. 44 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0034] FIGS. 45 and 46 are diagrammatic views that depict an embodiment of a system for data collection in accordance with the present disclosure.

[0035] FIGS. 47 and 48 are diagrammatic views that depict an embodiment of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0036] FIG. 49 depicts an embodiment of a data monitoring device incorporating sensors in accordance with the present disclosure.

[0037] FIGS. 50 and 51 are diagrammatic views that depict embodiments of a data monitoring device in communication with external sensors in accordance with the present disclosure.

[0038] FIG. 52 is a diagrammatic view that depicts embodiments of a data monitoring device with additional detail in the signal evaluation circuit in accordance with the present disclosure.

[0039] FIG. 53 is a diagrammatic view that depicts embodiments of a data monitoring device with additional detail in the signal evaluation circuit in accordance with the present disclosure.

[0040] FIG. 54 is a diagrammatic view that depicts embodiments of a data monitoring device with additional detail in the signal evaluation circuit in accordance with the present disclosure.

[0041] FIG. 55 is a diagrammatic view that depicts embodiments of a system for data collection in accordance with the present disclosure.

[0042] FIG. 56 is a diagrammatic view that depicts embodiments of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0043] FIG. 57 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0044] FIGS. 58 and 59 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0045] FIGS. 60 and 61 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0046] FIGS. 62 and 63 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0047] FIGS. 64 and 65 are diagrammatic views that depict embodiments of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0048] FIG. 66 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0049] FIGS. 67 and 68 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0050] FIG. 69 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0051] FIG. 70 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0052] FIGS. 71 and 72 are diagrammatic views that depict embodiments of a system for data collection in accordance with the present disclosure.

[0053] FIGS. 73 and 74 are diagrammatic views that depict embodiments of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0054] FIG. 75 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0055] FIGS. 76 and 77 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0056] FIG. 78 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0057] FIGS. 79 and 80 are diagrammatic views that depict embodiments of a system for data collection in accordance with the present disclosure.

[0058] FIGS. 81 and 82 are diagrammatic views that depict embodiments of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0059] FIG. 83 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0060] FIGS. 84 and 85 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0061] FIG. 86 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0062] FIGS. 87 and 88 are diagrammatic views that depict embodiments of a system for data collection in accordance with the present disclosure.

[0063] FIGS. 89 and 90 are diagrammatic views that depict embodiments of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0064] FIG. 91 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0065] FIGS. 92 and 93 are diagrammatic views that depict embodiments of a data monitoring device in accordance with the present disclosure.

[0066] FIG. 94 is a diagrammatic view that depicts embodiments of a data monitoring device in accordance with the present disclosure.

[0067] FIGS. 95 and 96 are diagrammatic views that depict embodiments of a system for data collection in accordance with the present disclosure.

[0068] FIGS. 97 and 98 are diagrammatic views that depict embodiments of a system for data collection comprising a plurality of data monitoring devices in accordance with the present disclosure.

[0069] FIGS. 99, 100, and 101 are diagrammatic views of components and interactions of a data collection architecture involving a collector of route templates and the routing of data collectors in an industrial environment in accordance with the present disclosure.

[0070] FIG. 102 is a diagrammatic view that depicts a monitoring system that employs data collection bands in accordance with the present disclosure.

[0071] FIG. 103 is a diagrammatic view that depicts a system that employs vibration and other noise in predicting states and outcomes in accordance with the present disclosure.

[0072] FIG. 104 is a diagrammatic view that depicts a system for data collection in an industrial environment in accordance with the present disclosure.

[0073] FIG. 105 is a diagrammatic view that depicts an apparatus for data collection in an industrial environment in accordance with the present disclosure.

[0074] FIG. 106 is a schematic flow diagram of a procedure for data collection in an industrial environment in accordance with the present disclosure.

[0075] FIG. 107 is a diagrammatic view that depicts a system for data collection in an industrial environment in accordance with the present disclosure.

[0076] FIG. 108 is a diagrammatic view that depicts an apparatus for data collection in an industrial environment in accordance with the present disclosure.

[0077] FIG. 109 is a schematic flow diagram of a procedure for data collection in an industrial environment in accordance with the present disclosure.

[0078] FIG. 110 is a diagrammatic view that depicts industry-specific feedback in an industrial environment in accordance with the present disclosure.

[0079] FIG. 111 is a diagrammatic view that depicts an exemplary user interface for smart band configuration of a system for data collection in an industrial environment in accordance with the present disclosure.

[0080] FIG. 112 is a diagrammatic view that depicts a graphical approach 11300 for back-calculation in accordance with the present disclosure.

[0081] FIG. 113 is a diagrammatic view that depicts a wearable haptic user interface device for providing haptic stimuli to a user that is responsive to data collected in an industrial environment by a system adapted to collect data in the industrial environment in accordance with the present disclosure.

[0082] FIG. 114 is a diagrammatic view that depicts an augmented reality display of heat maps based on data collected in an industrial environment by a system adapted to collect data in the environment in accordance with the present disclosure.

[0083] FIG. 115 is a diagrammatic view that depicts an augmented reality display including real time data overlaying a view of an industrial environment in accordance with the present disclosure.

[0084] FIG. 116 is a diagrammatic view that depicts a user interface display and components of a neural net in a graphical user interface in accordance with the present disclosure.

[0085] FIG. 117 is a diagrammatic view of components and interactions of a data collection architecture involving swarming data collectors and sensor mesh protocol in an industrial environment in accordance with the present disclosure.

[0086] FIG. 118 is a diagrammatic view that depicts a data collection system according to some aspects of the present disclosure.

[0087] FIG. 119 is a diagrammatic view that depicts a system for self-organized, network-sensitive data collection in an industrial environment in accordance with the present disclosure.

[0088] FIG. 120 is a diagrammatic view that depicts an apparatus for self-organized, network-sensitive data collection in an industrial environment in accordance with the present disclosure.

[0089] FIG. 121 is a diagrammatic view that depicts an apparatus for self-organized, network-sensitive data collection in an industrial environment in accordance with the present disclosure.

[0090] FIG. 122 is a diagrammatic view that depicts an apparatus for self-organized, network-sensitive data collection in an industrial environment in accordance with the present disclosure.

[0091] FIG. 123 and FIG. 124 are diagrammatic views that depict embodiments of transmission conditions in accordance with the present disclosure.

[0092] FIG. 125 is a diagrammatic view that depicts embodiments of a sensor data transmission protocol in accordance with the present disclosure.

[0093] FIG. 126 and FIG. 127 are diagrammatic views that depict embodiments of benchmarking data in accordance with the present disclosure.

[0094] FIG. 128 is a diagrammatic view that depicts embodiments of a system for data collection and storage in an industrial environment in accordance with the present disclosure.

[0095] FIG. 129 is a diagrammatic view that depicts embodiments of an apparatus for self-organizing storage for data collection for an industrial system in accordance with the present disclosure.

[0096] FIG. 130 is a diagrammatic view that depicts embodiments of a storage time definition in accordance with the present disclosure.

[0097] FIG. 131 is a diagrammatic view that depicts embodiments of a data resolution description in accordance with the present disclosure.

[0098] FIG. 132 and FIG. 133 diagrammatic views of an apparatus for self-organizing network coding for data collection for an industrial system in accordance with the present disclosure.

[0099] FIG. 134 and FIG. 135 diagrammatic views of data marketplace interacting with data collection in an industrial system in accordance with the present disclosure.

[0100] FIG. 136 is a diagrammatic view that depicts a smart heating system as an element in a network for an industrial Internet of Things ecosystem in accordance with the present disclosure.

[0101] FIG. 137 is a diagrammatic view that depicts an architecture, its components and functional relationships for an industrial Internet of Things solution in accordance with the present disclosure.

[0102] FIG. 138 is a schematic illustrating an example of a sensor kit deployed in an industrial setting according to some embodiments of the present disclosure.

[0103] FIG. 139 is a schematic illustrating an example of a sensor kit network having a star network topology according to some embodiments of the present disclosure.

[0104] FIG. 140 is a schematic illustrating an example of a sensor kit network having a mesh network topology according to some embodiments of the present disclosure.

[0105] FIG. 141 is a schematic illustrating an example of a sensor kit network having a hierarchical network topology according to some embodiments of the present disclosure.

[0106] FIG. 142 is a schematic illustrating an example of a sensor according to some embodiments of the present disclosure.

[0107] FIG. 143 is a schematic illustrating an example schema of a reporting packet according to some embodiments of the present disclosure.

[0108] FIG. 144 is a schematic illustrating an example of an edge device of a sensor kit according to some embodiments of the present disclosure.

[0109] FIG. 145 is a schematic illustrating an example of a backend system that receives sensor data from sensor kits deployed in industrial settings according to some embodiments of the present disclosure.

[0110] FIG. 146 is a flow chart illustrating an example set of operations of a method for encoding sensor data captured by a sensor kit according to some embodiments of the present disclosure.

[0111] FIG. 147 is a flow chart illustrating an example set of operations of a method for decoding sensor data provided to a backend system by a sensor kit according to some embodiments of the present disclosure.

[0112] FIG. 148 is a flow chart illustrating an example set of operations of a method for encoding sensor data captured by a sensor kit using a media codec according to some embodiments of the present disclosure.

[0113] FIG. 149 is a flow chart illustrating an example set of operations of a method for decoding sensor data provided to a backend system by a sensor kit using a media codec according to some embodiments of the present disclosure.

[0114] FIG. 150 is a flow chart illustrating an example set of operations of a method for determining a transmission strategy and/or a storage strategy for sensor data collected by a sensor kit based on the sensor data, according to some embodiments of the present disclosure.

[0115] Figures 151-155 are schematics illustrating different configurations of sensor kits according to some embodiments of the present disclosure.

[0116] FIG. 156 is a flowchart illustrating an example set of operations of a method for monitoring industrial settings using an automatically configured backend system, according to some embodiments of the present disclosure.

[0117] FIG. 157 is a plan view of a manufacturing facility illustrating an exemplary implementation of a sensor kit including an edge device, according to some embodiments of the present disclosure.

[0118] FIG. 158 is a plan view of a surface portion of an underwater industrial facility illustrating an exemplary implementation of a sensor kit including an edge device, according to some embodiments of the present disclosure.

[0119] FIG. 159 is a plan view of an indoor agricultural facility illustrating an exemplary implementation of a sensor kit including an edge device, according to some embodiments of the present disclosure.

[0120] FIG. 160 is a schematic illustrating an example of a sensor kit in communication with a data handling platform according to some embodiments of the present disclosure.

[0121] FIGS. 161-164 are diagrammatic views that depict embodiments of a system for using one or more wearable devices for mobile data collection in accordance with the present disclosure.

[0122] FIGS. 165, 166, and 167 are diagrammatic views that depict embodiments of a system for using one or more mobile robots and/or mobile vehicles for mobile data collection in accordance with the present disclosure.

[0123] FIGS. 168-171 are diagrammatic views that depict embodiments of a system for using one or more handheld devices for mobile data collection in accordance with the present disclosure.

[0124] FIGS. 172, 173, and 174 are diagrammatic views that depict embodiments of a computer vision system in accordance with the present disclosure.

[0125] FIGS. 175 and 176 are diagrammatic views that depict embodiments of a deep learning system for training a computer vision system in accordance with the present disclosure.

[0126] FIG. 177 depicts a predictive maintenance eco system network architecture.

[0127] FIG. 178 depicts finding service workers using machine learning for the predictive maintenance eco-system of FIG. 177.

[0128] FIG. 179 depicts ordering parts and service in a predictive maintenance eco-system.

[0129] FIG. 180 depicts deployment of smart RFID elements in an industrial machine environment.

[0130] FIG. 181 depicts a generalized data structure for machine information in a smart RFID.

[0131] FIG. 182 depicts a block level diagram of the storage structure of a smart RFID.

[0132] FIG. 183 depicts an example of data stored in a smart RFID.

[0133] FIG. 184 depicts a flow diagram of a method for collecting information from a machine.

[0134] FIG. 185 depicts a flow diagram of a method for collecting data from a production environment.

[0135] FIG. 186 depicts an on-line maintenance management system with interfaces for data sources updating information in the on-line maintenance management system data storage.

[0136] FIG. 187 depicts a distributed ledger for predictive maintenance information with role-specific access thereof.

[0137] FIG. 188 depicts a process for capturing images of portions of an industrial machine.

[0138] FIG. 189 depicts a process that uses machine learning on images to recognize a likely internal structure of an industrial machine.

[0139] FIG. 190 depicts a knowledge graph of the predictive maintenance gathering information.

[0140] FIG. 191 depicts an artificial intelligence system generating service recommendations and the like based on predictive maintenance analysis.

[0141] FIG. 192 depicts a predictive maintenance timeline superimposed on a preventive maintenance timeline.

[0142] FIG. 193 depicts a block diagram of potential sources of diagnostic information.

[0143] FIG. 194 depicts a diagram of a process for rating vendors.

[0144] FIG. 195 depicts a diagram of a process for rating procedures.

[0145] FIG. 196 depicts a diagram of Blockchain applied to transactions of a predictive maintenance eco-system.

[0146] FIG. 197 depicts a transfer function that facilitates converting vibration data into severity units.

[0147] FIG. 198 depicts a table that facilitates mapping vibration data to severity units.

[0148] FIG. 199 depicts a composite frequency graph for conventional vibration assessment and severity unit-based assessment.

[0149] FIG. 200 depicts a rendering of a portion of an industrial machine for use in an electronic user interface for depicting and discovering severity units and related information about a rotating component of the industrial machine.

[0150] FIG. 201 depicts a data table of rotating component design parameters for use in predicting maintenance events.

[0151] FIG. 202 is a flow chart of predicting maintenance of at least one of a gear, motor and roller bearing based on severity unit and actuator count, such as count of teeth in a gear.

[0152] FIG. 203 is a schematic diagram of an example platform for facilitating development of intelligence in an Industrial Internet of Things (IIoT) system according to some aspects of the present disclosure.

[0153] FIG. 204 is a schematic diagram showing additional details, components, sub-systems, and other elements of an optional implementation of the example platform of FIG. 203;

[0154] FIG. 205 is a schematic diagram showing a robotic process automation ("RPA") system of the example platform of FIG. 203;

[0155] FIG. 206 is a schematic diagram showing an opportunity mining system and an adaptive intelligence layer of the example platform of FIG. 203;

[0156] FIG. 207 is a schematic diagram showing optional elements of the adaptive intelligent systems layer that facilitate improved edge intelligence of the example platform of FIG. 203;

[0157] FIG. 208 is a schematic diagram showing optional elements of an industrial entity-oriented data storage systems layer of the example platform of FIG. 203;

[0158] FIG. 209 is a schematic diagram showing an example Robotic Process Automation system of the example platform of FIG. 203;

[0159] FIG. 210 is a schematic diagram of an example system for data processing in an industrial environment that utilizes protocol adaptors according to some aspects of the present disclosure;

[0160] FIG. 211 is another schematic diagram illustrating further components and elements of the example system of FIG. 210; and

[0161] FIG. 212 illustrates an example connect attempt of the example system of FIG. 210 according to some aspects of the present disclosure.

[0162] FIG. 213 is a schematic illustrating examples of architecture of a digital twin system according to embodiments of the present disclosure.

[0163] FIG. 214 is a schematic illustrating exemplary components of a digital twin management system according to embodiments of the present disclosure.

[0164] FIG. 215 is a schematic illustrating examples of a digital twin I/O system that interfaces with an environment, the digital twin system, and/or components thereof to provide bi-directional transfer of data between coupled components according to embodiments of the present disclosure.

[0165] FIG. 216 is a schematic illustrating examples of sets of identified states related to industrial environments that the digital twin system may identify and/or store for access by intelligent systems (e.g., a cognitive intelligence system) or users of the digital twin system according to embodiments of the present disclosure.

[0166] FIG. 217 is a schematic illustrating example embodiments of methods for updating a set of properties of a digital twin of the present disclosure on behalf of a client application and/or one or more embedded digital twins according to embodiments of the present disclosure.

[0167] FIG. 218 is a view of a display illustrating example embodiments of a display interface of the present disclosure that renders a digital twin of a dryer centrifuge with information relating to the dryer centrifuge according to embodiments of the present disclosure.

[0168] FIG. 219 is a schematic illustrating example embodiments of methods for updating a set of vibration fault level states of machine components such as bearings in the digital twin of an industrial machine, on behalf of a client application according to embodiments of the present disclosure.

[0169] FIG. 220 is a schematic illustrating example embodiments of methods for updating a set of vibration severity unit values of machine components such as bearings in the digital twin of a machine on behalf of a client application according to embodiments of the present disclosure.

[0170] FIG. 221 is a schematic illustrating example embodiments of a method for updating a set of probability of failure values in the digital twins of machine components on behalf of a client application according to embodiments of the present disclosure.

[0171] FIG. 222 is a schematic illustrating example embodiments of methods for updating a set of probability of downtime values of machines in the digital twin of a manufacturing facility on behalf of a client application according to embodiments of the present disclosure.

[0172] FIG. 223 is a schematic illustrating example embodiments of methods for updating a set of probability of shutdown values of manufacturing facilities in the digital twin of an enterprise on behalf of a client application according to embodiments of the present disclosure.

[0173] FIG. 224 is a schematic illustrating example embodiments of methods for updating a set of cost of downtime values of machines in the digital twin of a manufacturing facility according to embodiments of the present disclosure.

[0174] FIG. 225 is a schematic illustrating example embodiments of methods for updating one or more manufacturing KPI values in a digital twin of a manufacturing facility, on behalf of a client application according to embodiments of the present disclosure.

[0175] FIG. 226 is a view of a display illustrating further example embodiments of a display interface of the present disclosure that renders a digital twin of a dryer centrifuge with information relating to its drive components according to embodiments of the present disclosure.

[0176] FIG. 227 is a view of a display illustrating further example embodiments of a display interface of the present disclosure that provides a digital twin showing components of vibration according to embodiments of the present disclosure.

[0177] FIG. 228 is a view of a display illustrating further example embodiments of a display interface of the present disclosure that provides selections of digital twins showing various components experiencing faults according to embodiments of the present disclosure.

[0178] FIG. 229 is a view of a display illustrating example embodiments of a display interface of the present disclosure that renders a digital twin whose view incorporates connected machines each having drive bearings according to embodiments of the present disclosure.

[0179] FIG. 230 is a view of a display illustrating example embodiments of a display interface of the present disclosure that renders a digital twin whose view incorporates connected machines each having drive bearings showing motion outside of nominal according to embodiments of the present disclosure.

[0180] FIG. 231 is a view of a display illustrating example embodiments of a display interface of the present disclosure that renders a digital twin showing drive bearings corrected to nominal motion according to embodiments of the present disclosure.

[0181] FIG. 232 is a view of a display illustrating example embodiments of a display interface of the present disclosure that renders a digital twin whose view incorporates connected machines such as a motor and mill each having drive bearings showing motion outside of nominal according to embodiments of the present disclosure.

[0182] FIG. 233 is a view of a display illustrating example embodiments of a display interface of the present disclosure that renders a digital twin showing drive bearings corrected to nominal motion according to embodiments of the present disclosure.

[0183] FIG. 234 is a schematic illustrating an example of a portion of an information technology system for manufacturing artificial intelligence leveraging digital twins according to some embodiments of the present disclosure.

[0184] FIG. 235 is a schematic illustrating an example environment of the enterprise and industrial control tower and management platform, including data sources in communication therewith, according to some embodiments of the present disclosure.

[0185] FIG. 236 is a schematic illustrating an example implementation of the enterprise and industrial control tower and management platform according to some embodiments of the present disclosure.

[0186] FIG. 237 is a schematic illustrating an example set of components of the enterprise control tower and management platform according to some embodiments of the present disclosure.

[0187] FIG. 238 is a schematic illustrating an example of an enterprise data model according to some embodiments of the disclosure.

[0188] FIG. 239 is a schematic illustrating examples of different types of enterprise digital twins, including executive digital twins, in relation to the data layer, processing layer, and application layer of the enterprise digital twin framework according to some embodiments of the present disclosure.

[0189] FIG. 240 is a flow chart illustrating an example set of operations for configuring and serving an enterprise digital twin.

[0190] FIG. 241 is a schematic illustrating example embodiments of systems for fault diagnosis in an industrial environment having components according to embodiments of the disclosure.

[0191] FIG. 242 is a schematic illustrating example embodiments of methods for fault diagnosis in an industrial environment having components according to embodiments of the disclosure.

[0192] FIGS. 243-248 are views depicting implementations of the systems and the methods of the disclosure for fault diagnosis in an industrial environment having components according to example embodiments of the disclosure.

[0193] FIGS. 249-252 are schematics illustrating example embodiments of architectures for implementation of the systems and the methods of the disclosure for fault diagnosis in an industrial environment having components according to embodiments of the disclosure.

[0194] FIG. 253 is a schematic illustrating an example of a portion of an information technology system for manufacturing artificial intelligence leveraging digital twins according to some embodiments of the disclosure.

[0195] FIG. 254 is a schematic view of an exemplary embodiment of the quantum computing service according to some embodiments of the present disclosure.

[0196] FIG. 255 illustrates quantum computing service request handling according to some embodiments of the present disclosure.

[0197] FIG. 256 is a diagrammatic view that illustrates embodiments of the biology-based industrial internet of things system in accordance with the present disclosure.

[0198] FIG. 257 is a diagrammatic view of the thalamus service and how it coordinates within the modules in accordance with the present disclosure.

[0199] FIG. 258 is a diagrammatic view of a dual process artificial neural network system in accordance with the present disclosure.

[0200] FIG. 259 is a diagrammatic view illustrating an example implementation of a conventional computer vision system for creating an image of an object of interest.

[0201] FIG. 260 is a schematic illustrating an example implementation of a dynamic vision system for dynamically learning an object concept about an object of interest according to some embodiments of the present disclosure.

[0202] FIG. 261 is a schematic illustrating an example architecture of a dynamic vision system according to some embodiments of the present disclosure.

[0203] FIG. 262 is a flow diagram illustrating a method for object recognition by a dynamic vision system according to some embodiments of the present disclosure.

[0204] FIG. 263 is a schematic illustrating an example implementation of a dynamic vision system for modelling, simulating and optimizing various optical, mechanical, design and lighting parameters of the dynamic vision system according to some embodiments of the present disclosure.

[0205] FIG. 264 is a schematic illustrating an example artificial neural network used to provide real-time, adaptive control of a dynamic vision system according to some embodiments of the present disclosure.

[0206] FIG. 265 is a diagrammatic view illustrating an example implementation of a dynamic vision system using a convolutional neural network (CNN) to provide classification of an object of interest according to some embodiments of the present disclosure.

[0207] FIG. 266 is a diagrammatic view illustrating an example implementation of a dynamic vision system using a transformer network to provide classification of an object of interest according to some embodiments of the present disclosure.

[0208] FIG. 267 is a schematic view illustrating an example implementation of a dynamic vision system depicting detailed view of various components along with integration of the dynamic vision system with one or more third party systems according to some embodiments of the present disclosure.

[0209] FIG. 268 is a block diagram depicting exemplary features, capabilities, and interfaces of a robust generative artificial intelligence platform.

[0210] FIG. 269 is a diagram illustrating a determination of attention by a machine learning model.

[0211] FIG. 270 is a diagram of a transformer model.

[0212] FIG. 271 is a schematic view of an example AI convergence system of systems.

[0213] FIG. 272 is a schematic view of an example offering layer.

[0214] FIG. 273 is a schematic view of an example transactions layer.

[0215] FIG. 274 is a schematic view of an example operations layer.

[0216] FIG. 275 is a schematic view of an example network layer.

[0217] FIG. 276 is a schematic view of an example data layer.

[0218] FIG. 277 is a schematic view of an example data layer.

[0219] FIG. 278 is a schematic view of an example intelligent data layer architecture.

[0220] FIG. 279 is a schematic view of an example network layer.

[0221] FIG. 280 is a schematic view of an example AI subsystem integrator system.

[0222] FIG. 281 is a schematic view of an example multiplatform attention management.

[0223] FIG. 282 is a diagram of an industrial technology stack in accordance with some example embodiments.

[0224] FIG. 283 is a diagram of a set of examples of industrial modules of the industrial technology stack in accordance with some example embodiments.

DETAILED DESCRIPTION

[0225] Detailed embodiments of the present disclosure are disclosed herein; however, it is to be understood that the disclosed embodiments are merely exemplary of the disclosure, which may be embodied in various forms. Therefore, specific structural and functional details disclosed herein are not to be interpreted as limiting, but merely as a basis for the claims and as a representative

basis for teaching one skilled in the art to variously employ the present disclosure in virtually any appropriately detailed structure.

ARTIFICIAL INTELLIGENCE SYSTEMS AND TECHNIQUES FOR INDUSTRIAL INTERNET OF THINGS

Computer Vision and Neural Network Systems for Industrial Internet of Things

[0226] In example embodiments, the disclosure provides a computer-implemented method for fault diagnosis in an industrial environment having a plurality of components. The computer-implemented method includes providing a plurality of sensors to the industrial environment, each of the plurality of sensors may be operatively coupled to at least one of the plurality of components and configured to generate a plurality of sensor data values in response to one or more sensed parameters. The plurality of sensor data values may be processed to determine a recognized pattern therefrom. At least one industrial-environment digital twin corresponding to the industrial environment may be retrieved. The at least one industrial-environment digital twin may include a plurality of component digital twins, with each of the plurality of component digital twins corresponding to one of the plurality of components in the industrial environment. The at least one industrial-environment digital twin and the plurality of component digital twins may be visual digital twins that may be configured to be rendered in a visual manner. The at least one industrial-environment digital twin and at least one respective component digital twin of the plurality of component digital twins may be updated based on the plurality of sensor data values, at least in part, in response to the determination of the recognized pattern for the corresponding component. A request may be received from a client application to check an operational condition of a particular component from the plurality of components in the industrial environment. The at least one industrial-environment digital twin and the at least one respective component digital twin corresponding to the particular component in the client application may be rendered in response to the received request and based on the operational condition of the particular component.

[0227] In example embodiments, the computer-implemented method may further include determining if the recognized pattern relates to at least one system characteristic including at least one of: a fault operation for a given component of the plurality of components, an off-nominal operation for the given component of the plurality of components, or an exceedance value for the given component of the plurality of components. Such the at least one system characteristic may be generally indicative of some fault in the corresponding component in the industrial environment and thus may be useful for the purposes of the disclosure.

[0228] In example embodiments, the computer-implemented method may further include generating a notification in the client application in response to a determination that the recognized pattern relates to the at least one system characteristic for the given component. In example embodiments, the computer-implemented method may further comprise configuring the client application to allow selection of the notification. The rendering the at least one industrial-environment digital twin and the at least one respective component digital twin corresponding to the given component may be in response to the selection of the notification. Such client application may be installed on a client device and allows the client to conveniently access information related to any fault determination in the industrial environment.

[0229] In example embodiments, the rendering may further comprise executing a simulation for the at least one industrial-environment digital twin and the at least one respective component digital twin based on the recognized pattern. In example embodiments, the simulation may simulate an effect of the recognized pattern on an operation of the corresponding component. In example embodiments, the rendering may further comprise executing another second simulation for the at least one industrial-environment digital twin and the at least one respective component digital twin based on a normal operation of the corresponding component. With such a rendering, the client may be provided with sufficient visual information to diagnose the fault in the industrial environment.

[0230] In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin to the client application may be via a display device of a user device. In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin to the client application may be via an augmented reality-enabled device. In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin to the client application may be via a virtual reality headset.

[0231] In example embodiments, the plurality of sensors may comprise at least one vibration measurement sensor coupled to a motor of the corresponding component. The one or more sensed parameters may comprise vibration parameters related to a wobble in the motor of the corresponding component. In example embodiments, the recognized pattern may comprise at least one of: a broken bearing in the motor, broken or cracked rotor bars in the motor, a misalignment in the motor, an imbalance in the motor, or a material build-up in the motor. In example embodiments, the one or more sensed parameters may include at least one of: a set of temperature parameters, pressure parameters, humidity parameters, wind parameters, rainfall parameters, tide parameters, storm surge parameters, cloud cover parameters, snowfall parameters, visibility parameters, radiation parameters, audio parameters, video parameters, image parameters, water level parameters, quantum parameters, flow rate parameters, signal power parameters, signal frequency parameters, motion parameters, velocity parameters, acceleration parameters, lighting level parameters, analyte concentration parameters, biological compound concentration parameters, metal concentration parameters, or organic compound concentration parameters.

[0232] In example embodiments, the plurality of component digital twins may be generated based on properties of the corresponding component imported from at least one of: respective manufacturers of the components, onboard libraries, crowdsourced material, or subscription marketplaces.

[0233] In example embodiments, the computer-implemented method may further comprise providing an executive digital twin configured to provide forecasted financial information for the given component based, at least in part, on the at least one system characteristic determined to be related to the recognized pattern. In example embodiments, the computer-implemented method may further comprise providing an operator digital twin configured to provide workflow

information for performing maintenance for the given component based, at least in part, on the at least one system characteristic determined to be related to the recognized pattern.

[0234] In example embodiments, the rendering the at least one industrial-environment digital twin may include rendering the at least one industrial-environment digital twin as a digital representation of a real world element. In example embodiments, the rendering the at least one industrial-environment digital twin may include at least one of mimicking, copying, or modeling behaviors of the real world element in response to at least one of inputs, outputs, or conditions of an environment. In example embodiments, the rendering the at least one respective component digital twin corresponding to the particular component may include rendering the at least one respective component digital twins as a set of discrete component digital twins embedded within the at least one industrial-environment digital twin. In example embodiments, the rendering the set of discrete component digital twins may include rendering the set of discrete component digital twins based on imported properties of the particular component and on historical behavior of the particular component for implementation in the industrial environment.

[0235] In example embodiments, the method may further include providing an operator digital twin configured to generate visual cues indicating potential problems with an identified component of the plurality of components. In example embodiments, the providing the operator digital twin may further include generating a selector for selection by a user to direct maintenance on the identified component and the method may further include directing the maintenance on the identified component in response to selection of the selector.

[0236] In example embodiments, the method may further include generating at least one of a picture or a video of a component in response to an instruction from a user and further including detecting wobble induced by bad poles based on the at least one of the picture or the video. In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin may be in response to selection of a received request.

[0237] In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin may include rendering the at least one industrial-environment digital twin and the at least one respective component digital twin in a visual manner. The method may further include drilling down on a particular element to view additional information regarding the particular element in response to a selection by a user on a display corresponding to the at least one industrial-environment digital twin and the at least one respective component digital twin as rendered in the visual manner.

[0238] In example embodiments, the disclosure provides a computing system for fault diagnosis in an industrial environment having a plurality of components. The computing system may comprise a plurality of sensors associated with the industrial environment, with each of the plurality of sensors operatively coupled to at least one of the plurality of components. The plurality of sensors may be configured to generate a plurality of sensor data values in response to one or more sensed parameters. At least one industrial-environment digital twin may correspond to the industrial environment. The at least one industrial-environment digital twin may include a plurality of component digital twins, with each of the plurality of component digital twins corresponding to

one of the plurality of components in the industrial environment. The at least one industrial-environment digital twin and the plurality of component digital twins may be visual digital twins that may be configured to be rendered in a visual manner. One or more processors may be configured to: process the plurality of sensor data values to determine a recognized pattern therefrom; update the at least one industrial-environment digital twin and at least one respective component digital twin of the plurality of component digital twins based on the plurality of sensor data values, at least in part, in response to the determination of the recognized pattern for the corresponding component; receive a request from a client application to check an operational condition of a particular component from the plurality of components in the industrial environment; and render the at least one industrial-environment digital twin and the at least one respective component digital twin corresponding to the particular component in the client application in response to the received request and based on the operational condition of the particular component.

[0239] In example embodiments, the system may further comprise an executive digital twin configured to provide forecasted financial information for a given component based, at least in part, on at least one system characteristic determined to be related to the recognized pattern. In example embodiments, the system may further comprise an operator digital twin configured to provide workflow information for performing maintenance for a given component based, at least in part, on at least one system characteristic determined to be related to the recognized pattern.

[0240] In example embodiments, the one or more processors may be further configured to determine if the recognized pattern relates to at least one system characteristic including at least one of: a fault operation for a given component of the plurality of components, an off-nominal operation for the given component of the plurality of components, or an exceedance value for the given component of the plurality of components. In example embodiments, the one or more processors may be further configured to generate a notification in the client application in response to the determination that the recognized pattern may relate to the at least one system characteristic for the given component. In example embodiments, the one or more processors may be further configured to configure the client application to allow selection of the notification, and where the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin may correspond to the given component is in response to the selection of the notification. In example embodiments, the plurality of sensors may be configured to generate the plurality of sensor data values to include a stream of phase-based data for at least one of temperature, humidity, or load. In example embodiments, the plurality of sensors may be configured to generate at least one of a continuous stream of data over time, a nearly continuous stream of data over time, periodic readings, event-driven readings, or readings according to a selected schedule. In example embodiments, the plurality of sensor data values may include vibration parameters related to a wobble in a motor of the at least one of the plurality of components, and where the one or more processors may be further configured to generate maintenance indications based on the vibration parameters related to the wobble. In example embodiments, the one or more processors may be further configured to at least one of: predict a

bearing life for the motor, identify a bearing health parameter, identify a bearing performance parameter, identify wear on a bearing, identify presence of foreign matter in bearings, identify air gaps in bearings, identify a loss of fluid in fluid coated bearings, identify stress and strain of flexure bearings, or identify behavior at a selected operation frequency for the plurality of components.

[0241] In example embodiments, the disclosure may provide a non-transitory computer readable storage medium having a plurality of instructions stored thereon which, when executed across one or more processors, causes at least a portion of the one or more processors to perform operations comprising: providing a plurality of sensors to an industrial environment having a plurality of components, where each of the plurality of sensors operatively may be coupled to at least one of the plurality of components and configured to generate a plurality of sensor data values in response to one or more sensed parameters; processing the plurality of sensor data values to determine a recognized pattern therefrom; retrieving at least one industrial-environment digital twin corresponding to the industrial environment, where the at least one industrial-environment digital twin may include a plurality of component digital twins, with each of the plurality of component digital twins corresponding to one of the plurality of components in the industrial environment, and where the at least one industrial-environment digital twin and the plurality of component digital twins may be visual digital twins that may be configured to be rendered in a visual manner; updating the at least one industrial-environment digital twin and at least one respective component digital twin of the plurality of component digital twins based on the plurality of sensor data values, at least in part, in response to determination of the recognized pattern for the corresponding component; receiving a request from a client application to check an operational condition of a particular component from the plurality of components in the industrial environment; and rendering the at least one industrial-environment digital twin and the at least one respective component digital twin corresponding to the particular component in the client application in response to the received request and based on the operational condition of the particular component.

[0242] In example embodiments, a maintenance system for an industrial environment may include a plurality of industrial machines, a predictive maintenance system, and a maintenance notification system. The plurality of industrial machines may collectively include a plurality of motors, the plurality of motors collectively including a predefined number of rotor bars. The predictive maintenance system may be programmed to generate a maintenance schedule for the plurality of industrial machines based on the predefined number of rotor bars and a rotor bar failure rate formula. The maintenance notification system may be programmed to generate maintenance alerts to indicate that maintenance should be performed on the plurality of industrial machines based on the maintenance schedule. In example embodiments, the rotor bar failure rate formula may be based on rotor bar weakening. In example embodiments, each of the plurality of motors may have a cycle rate and an age, and the predictive maintenance system may be further programmed to generate the maintenance schedule based on the cycle rate and the age of each of the plurality of motors.

[0243] One aspect of the current disclosure relates to a method for transmitting a predictive model of a data stream from a first device to a second device. The method may include receiving, by a first device, a plurality of data values of a data stream. The data values may comprise sensor data collected from one or more sensor devices. The method may include generating, by the first device, a predictive model for predicting future data values of the data stream based on the received plurality of data values. Generating the predictive model may include determining a plurality of model parameters. The method may include transmitting, by the first device, the plurality of model parameters to the second device. The method may include receiving, by the second device, the plurality of model parameters. The method may include parameterizing, by the second device, a predictive model using the plurality of model parameters. The method may include predicting, by the second device, the future data values of the data stream using the parameterized predictive model. In example embodiments, the parameters comprise a vector. In example embodiments, the vector is a motion vector associated with a robot. In example embodiments, the future data values of the data stream comprise one or more future predicted locations of the robot. In example embodiments, the predictive model predicts stock levels of items, the method further including detecting, based on the future data values, an upcoming supply shortage of an item. The method may further include taking action to avoid running out of the item. In example embodiments, the predictive model is a behavior analysis model. In example embodiments, the future data values indicate a predicted behavior of an entity. In example embodiments, the predictive model is an augmentation model, wherein the future data values correspond to an inoperative sensor. In example embodiments, the predictive model is a classification model. In example embodiments, the future data values indicate a predicted future state of a system comprising the one or more sensor devices. In example embodiments, the sensors are security cameras. In example embodiments, the data stream comprises motion vectors extracted from video data captured by the security cameras. In example embodiments, the sensors are vibration sensors measuring vibrations generated by machines. In example embodiments, the future data values indicate a potential need for maintenance of the machines. The method may further include receiving, by the first device, additional data values of the data stream. The method may include refining, by the first device, the predictive model using the additional data values. In example embodiments, refining the predictive model adjusts the model parameters. The method may include transmitting the adjusted model parameters to the second device. The method may further include receiving, by the second device, the adjusted model parameters. The method may include re-parameterizing the predictive model using the adjusted model parameters. The method may include generating additional future data values using the re-parameterized predictive model.

[0244] Another aspect of the current disclosure relates to a method for prioritizing predictive model data streams. The method may include receiving, by a first device, a plurality of predictive model data streams. In example embodiments, each predictive model data streams comprises a set of model parameters for a corresponding predictive model. In example embodiments, each predictive model is trained to predict future data values of a data source. The method may include prioritizing, by the first device, priorities to each of the plurality of predictive model data streams.

The method may include selecting at least one of the predictive model data streams based on a corresponding priority. The method may include parameterizing, by the first device, a predictive model using the set of model parameters included in the selected predictive model stream. The method may include predicting, by the first device, future data values of the data source using the parameterized predictive model. In example embodiments, the selected at least one predictive model data stream is associated with a high priority. In example embodiments, the selecting comprises suppressing the predictive model data streams that were not selected based on the priorities associated with each non-selected predictive model data stream. In example embodiments, assigning priorities to each of the plurality of predictive model data streams comprises determining whether each set of model parameters is unusual. In example embodiments, assigning priorities to each of the plurality of predictive model data streams comprises determining whether each set of model parameters has changed from a previous value. In example embodiments, the set of model parameters comprise at least one vector. In example embodiments, the at least one vector comprises a motion vector associated with a robot. In example embodiments, the future data values comprise one or more future predicted locations of the robot. In example embodiments, the predictive model is a behavior analysis model. In example embodiments, future data values indicate a predicted behavior of an entity. In example embodiments, the predictive model is an augmentation model. In example embodiments, the future data values correspond to an inoperative sensor. In example embodiments, the predictive model is a classification model. In example embodiments, the future data values indicate a predicted future state of a system comprising the one or more sensor devices. In example embodiments, the sensors are security cameras. In example embodiments, the data stream comprises motion vectors extracted from video data captured by the security cameras. In example embodiments, the sensors are vibration sensors measuring vibrations generated by machines. In example embodiments, the future data values indicate a potential need for maintenance of the machines.

[0245] Methods and systems described herein for industrial machine sensor data streaming, collection, processing, and storage may be configured to operate with existing data collection, processing, and storage systems while preserving access to existing format/frequency range/resolution compatible data. While the industrial machine sensor data streaming facilities described herein may collect a greater volume of data (e.g., longer duration of data collection) from sensors at a wider range of frequencies and with greater resolution than existing data collection systems, methods and systems may be employed to provide access to data from the stream of data that represents one or more ranges of frequency and/or one or more lines of resolution that are purposely compatible with existing systems. Further, a portion of the streamed data may be identified, extracted, stored, and/or forwarded to existing data processing systems to facilitate operation of existing data processing systems that substantively matches operation of existing data processing systems using existing collection-based data. In this way, a newly deployed system for sensing aspects of industrial machines, such as aspects of moving parts of industrial machines, may facilitate continued use of existing sensed data processing facilities, algorithms, models, pattern recognizers, user interfaces, and the like.

[0246] Through identification of existing frequency ranges, formats, and/or resolution, such as by accessing a data structure that defines these aspects of existing data, higher resolution streamed data may be configured to represent a specific frequency, frequency range, format, and/or resolution. This configured streamed data can be stored in a data structure that is compatible with existing sensed data structures so that existing processing systems and facilities can access and process the data substantially as if it were the existing data. One approach to adapting streamed data for compatibility with existing sensed data may include aligning the streamed data with existing data so that portions of the streamed data that align with the existing data can be extracted, stored, and made available for processing with existing data processing methods. Alternatively, data processing methods may be configured to process portions of the streamed data that correspond, such as through alignment, to the existing data, with methods that implement functions substantially similar to the methods used to process existing data, such as methods that process data that contain a particular frequency range or a particular resolution and the like.

[0247] Methods used to process existing data may be associated with certain characteristics of sensed data, such as certain frequency ranges, sources of data, and the like. As an example, methods for processing bearing sensing information for a moving part of an industrial machine may be capable of processing data from bearing sensors that fall into a particular frequency range. This method can thusly be at least partially identifiable by these characteristics of the data being processed. Therefore, given a set of conditions, such as moving device being sensed, industrial machine type, frequency of data being sensed, and the like, a data processing system may select an appropriate method. Also, given such a set of conditions, an industrial machine data sensing and processing facility may configure elements, such as data filters, routers, processors, and the like, to handle data meeting the conditions.

[0248] Figures 1 through 5 depict portions of an overall view of an industrial Internet of Things (IoT) data collection, monitoring and control system 10. Figure 2 depicts a mobile ad hoc network ("MANET") 20, which may form a secure, temporal network connection 22 (sometimes connected and sometimes isolated), with a cloud 30 or other remote networking system, so that network functions may occur over the MANET 20 within the environment, without the need for external networks, but at other times information can be sent to and from a central location. This allows the industrial environment to use the benefits of networking and control technologies, while also providing security, such as preventing cyber-attacks. The MANET 20 may use cognitive radio technologies 40, including those that form up an equivalent to the IP protocol, such as router 42, MAC 44, and physical layer technologies 46. In embodiments, the system depicted in Figures 1 through 5 provides network-sensitive or network-aware transport of data over the network to and from a data collection device or a heavy industrial machine.

[0249] Figures 3-4 depict intelligent data collection technologies deployed locally, at the edge of an IoT deployment, where heavy industrial machines are located. This includes various sensors 52, IoT devices 54, data storage capabilities (e.g., data pools 60, or distributed ledger 62) (including intelligent, self-organizing storage), sensor fusion (including self-organizing sensor fusion), and the like. Interfaces for data collection, including multi-sensory interfaces, tablets, smartphones 58,

and the like are shown. Figure 3 also shows data pools 60 that may collect data published by machines or sensors that detect conditions of machines, such as for later consumption by local or remote intelligence. A distributed ledger system 62 may distribute storage across the local storage of various elements of the environment, or more broadly throughout the system. Figure 4 also shows on-device sensor fusion 80, such as for storing on a device data from multiple analog sensors 82, which may be analyzed locally or in the cloud, such as by machine learning 84, including by training a machine based on initial models created by humans that are augmented by providing feedback (such as based on measures of success) when operating the methods and systems disclosed herein.

[0250] Figure 1 depicts a server based portion of an industrial IoT system that may be deployed in the cloud or on an enterprise owner's or operator's premises. The server portion includes network coding (including self-organizing network coding and/or automated configuration) that may configure a network coding model based on feedback measures, network conditions, or the like, for highly efficient transport of large amounts of data across the network to and from data collection systems and the cloud. Network coding may provide a wide range of capabilities for intelligence, analytics, remote control, remote operation, remote optimization, various storage configurations and the like, as depicted in Figure 1. The various storage configurations may include distributed ledger storage for supporting transactional data or other elements of the system.

[0251] Figure 5 depicts a programmatic data marketplace 70, which may be a self-organizing marketplace, such as for making available data that is collected in industrial environments, such as from data collectors, data pools, distributed ledgers, and other elements disclosed herein. Additional detail on the various components and sub-components of Figures 1 through 5 is provided throughout this disclosure.

[0252] With reference to Figure 6, an embodiment of platform 100 may include a local data collection system 102, which may be disposed in an environment 104, such as an industrial environment similar to that shown in Figure 3, for collecting data from or about the elements of the environment, such as machines, components, systems, sub-systems, ambient conditions, states, workflows, processes, and other elements. The platform 100 may connect to or include portions of the industrial IoT data collection, monitoring and control system 10 depicted in Figures 1-5. The platform 100 may include a network data transport system 108, such as for transporting data to and from the local data collection system 102 over a network 110, such as to a host processing system 112, such as one that is disposed in a cloud computing environment or on the premises of an enterprise, or that consists of distributed components that interact with each other to process data collected by the local data collection system 102. The host processing system 112, referred to for convenience in some cases as the host system 112, may include various systems, components, methods, processes, facilities, and the like for enabling automated, or automation-assisted processing of the data, such as for monitoring one or more environments 104 or networks 110 or for remotely controlling one or more elements in a local environment 104 or in a network 110. The platform 100 may include one or more local autonomous systems, such as for enabling autonomous behavior, such as reflecting artificial, or machine-based intelligence or such as enabling automated

action based on the applications of a set of rules or models upon input data from the local data collection system 102 or from one or more input sources 116, which may comprise information feeds and inputs from a wide array of sources, including those in the local environment 104, in a network 110, in the host system 112, or in one or more external systems, databases, or the like. The platform 100 may include one or more intelligent systems 118, which may be disposed in, integrated with, or acting as inputs to one or more components of the platform 100. Details of these and other components of the platform 100 are provided throughout this disclosure.

[0253] Intelligent systems 118 may include cognitive systems 120, such as enabling a degree of cognitive behavior as a result of the coordination of processing elements, such as mesh, peer-to-peer, ring, serial, and other architectures, where one or more node elements is coordinated with other node elements to provide collective, coordinated behavior to assist in processing, communication, data collection, or the like. The MANET 20 depicted in Figure 2 may also use cognitive radio technologies, including those that form up an equivalent to the IP protocol, such as router 42, MAC 44, and physical layer technologies 46. In one example, the cognitive system technology stack can include examples disclosed in U.S. Patent Number 8,060,017 to Schlicht et al., issued 15 November 2011 and hereby incorporated by reference as if fully set forth herein.

[0254] Intelligent systems may include machine learning systems 122, such as for learning on one or more data sets. The one or more data sets may include information collected using local data collection systems 102 or other information from input sources 116, such as to recognize states, objects, events, patterns, conditions, or the like that may, in turn, be used for processing by the host system 112 as inputs to components of the platform 100 and portions of the industrial IoT data collection, monitoring and control system 10, or the like. Learning may be human-supervised or fully-automated, such as using one or more input sources 116 to provide a data set, along with information about the item to be learned. Machine learning may use one or more models, rules, semantic understandings, workflows, or other structured or semi-structured understanding of the world, such as for automated optimization of control of a system or process based on feedback or feed forward to an operating model for the system or process. One such machine learning technique for semantic and contextual understandings, workflows, or other structured or semi-structured understandings is disclosed in U.S. Patent Number 8,200,775 to Moore, issued 12 June 2012, and hereby incorporated by reference as if fully set forth herein. Machine learning may be used to improve the foregoing, such as by adjusting one or more weights, structures, rules, or the like (such as changing a function within a model) based on feedback (such as regarding the success of a model in a given situation) or based on iteration (such as in a recursive process). Where sufficient understanding of the underlying structure or behavior of a system is not known, insufficient data is not available, or in other cases where preferred for various reasons, machine learning may also be undertaken in the absence of an underlying model; that is, input sources may be weighted, structured, or the like within a machine learning facility without regard to any a priori understanding of structure, and outcomes (such as those based on measures of success at accomplishing various desired objectives) can be serially fed to the machine learning system to allow it to learn how to achieve the targeted objectives. For example, the system may learn to

recognize faults, to recognize patterns, to develop models or functions, to develop rules, to optimize performance, to minimize failure rates, to optimize profits, to optimize resource utilization, to optimize flow (such as flow of traffic), or to optimize many other parameters that may be relevant to successful outcomes (such as outcomes in a wide range of environments). Machine learning may use genetic programming techniques, such as promoting or demoting one or more input sources, structures, data types, objects, weights, nodes, links, or other factors based on feedback (such that successful elements emerge over a series of generations). For example, alternative available sensor inputs for a data collection system 102 may be arranged in alternative configurations and permutations, such that the system may, using generic programming techniques over a series of data collection events, determine what permutations provide successful outcomes based on various conditions (such as conditions of components of the platform 100, conditions of the network 110, conditions of a data collection system 102, conditions of an environment 104), or the like. In example embodiments, local machine learning may turn on or off one or more sensors in a multi-sensor data collector 102 in permutations over time, while tracking success outcomes such as contributing to success in predicting a failure, contributing to a performance indicator (such as efficiency, effectiveness, return on investment, yield, or the like), contributing to optimization of one or more parameters, identification of a pattern (such as relating to a threat, a failure mode, a success mode, or the like) or the like. For example, a system may learn what sets of sensors should be turned on or off under given conditions to achieve the highest value utilization of a data collector 102. In example embodiments, similar techniques may be used to handle optimization of transport of data in the platform 100 (such as in the network 110) by using generic programming or other machine learning techniques to learn to configure network elements (such as configuring network transport paths, configuring network coding types and architectures, configuring network security elements), and the like.

[0255] In example embodiments, the local data collection system 102 may include a high-performance, multi-sensor data collector having a number of novel features for collection and processing of analog and other sensor data. In example embodiments, a local data collection system 102 may be deployed to the industrial facilities depicted in Figure 3. A local data collection system 102 may also be deployed to monitor other machines such as the machine 2200. The data collection system 102 may have on-board intelligent systems 118 (such as for learning to optimize the configuration and operation of the data collector, such as configuring permutations and combinations of sensors based on contexts and conditions). In one example, the data collection system 102 includes a crosspoint switch 130 or other analog switches. Automated, intelligent configuration of the local data collection system 102 may be based on a variety of types of information, such as information from various input sources, including those based on available power, power requirements of sensors, the value of the data collected (such as based on feedback information from other elements of the platform 100), the relative value of information (such as values based on the availability of other sources of the same or similar information), power availability (such as for powering sensors), network conditions, ambient conditions, operating states, operating contexts, operating events, and many others.

[0256] Figure 7 shows elements and sub-components of a data collection and analysis system 1100 for sensor data (such as analog sensor data) collected in industrial environments. As depicted in Figure 7, embodiments of the methods and systems disclosed herein may include hardware that has several different modules starting with the multiplexer ("MUX") main board 1104. In example embodiments, there may be a MUX option board 1108. The MUX 114 main board is where the sensors connect to the system. These connections are on top to enable ease of installation. Then there are numerous settings on the underside of this board as well as on the Mux option board 1108, which attaches to the MUX main board 1104 via two headers one at either end of the board. In example embodiments, the Mux option board has the male headers, which mesh together with the female header on the main Mux board. This enables them to be stacked on top of each other thus taking up less real estate.

[0257] In example embodiments, the main Mux board and/or the MUX option board then connects to the mother (e.g., with 4 simultaneous channels) and daughter (e.g., with 4 additional channels for 8 total channels) analog boards 1110 via cables where some of the signal conditioning (such as hardware integration) occurs. The signals then move from the analog boards 1110 to an anti-aliasing board (not shown) where some of the potential aliasing is removed. The rest of the aliasing removal is done on the delta sigma board 1112. The delta sigma board 1112 provides more aliasing protection along with other conditioning and digitizing of the signal. Next, the data moves to the Jennic™ board 1114 for more digitizing as well as communication to a computer via USB or Ethernet. In example embodiments, the Jennic™ board 1114 may be replaced with a pic board 1118 for more advanced and efficient data collection as well as communication. Once the data moves to the computer software 1102, the computer software 1102 can manipulate the data to show trending, spectra, waveform, statistics, and analytics.

[0258] In example embodiments, the system is meant to take in all types of data from volts to 4-20 mA signals. In example embodiments, open formats of data storage and communication may be used. In some instances, certain portions of the system may be proprietary especially some of research and data associated with the analytics and reporting. In example embodiments, smart band analysis is a way to break data down into easily analyzed parts that can be combined with other smart bands to make new more simplified yet sophisticated analytics. In example embodiments, this unique information is taken and graphics are used to depict the conditions because picture depictions are more helpful to the user. In example embodiments, complicated programs and user interfaces are simplified so that any user can manipulate the data like an expert.

[0259] In example embodiments, the system in essence, works in a big loop. The system starts in software with a general user interface ("GUI") 1124. In example embodiments, rapid route creation may take advantage of hierarchical templates. In example embodiments, a GUI is created so any general user can populate the information itself with simple templates. Once the templates are created the user can copy and paste whatever the user needs. In addition, users can develop their own templates for future ease of use and to institutionalize the knowledge. When the user has entered all of the user's information and connected all of the user's sensors, the user can then start the system acquiring data.

[0260] Embodiments of the methods and systems disclosed herein may include unique electrostatic protection for trigger and vibration inputs. In many critical industrial environments where large electrostatic forces, which can harm electrical equipment, may build up, for example rotating machinery or low-speed balancing using large belts, proper transducer and trigger input protection is required. In example embodiments, a low-cost but efficient method is described for such protection without the need for external supplemental devices.

[0261] Typically, vibration data collectors are not designed to handle large input voltages due to the expense and the fact that, more often than not, it is not needed. A need exists for these data collectors to acquire many varied types of RPM data as technology improves and monitoring costs plummet. In example embodiments, a method is using the already established OptoMOS™ technology which permits the switching up front of high voltage signals rather than using more conventional reed-relay approaches. Many historic concerns regarding non-linear zero crossing or other non-linear solid-state behaviors have been eliminated with regard to the passing through of weakly buffered analog signals. In addition, in example embodiments, printed circuit board routing topologies place all of the individual channel input circuitry as close to the input connector as possible. In example embodiments, a unique electrostatic protection for trigger and vibration inputs may be placed upfront on the Mux and DAQ hardware in order to dissipate the built up electric charge as the signal passed from the sensor to the hardware. In example embodiments, the Mux and analog board may support high-amperage input using a design topology comprising wider traces and solid state relays for upfront circuitry.

[0262] In some systems multiplexers are afterthoughts and the quality of the signal coming from the multiplexer is not considered. As a result of a poor quality multiplexer, the quality of the signal can drop as much as 30 dB or more. Thus, substantial signal quality may be lost using a 24-bit DAQ that has a signal to noise ratio of 110 dB and if the signal to noise ratio drops to 80 dB in the Mux, it may not be much better than a 16-bit system from 20 years ago. In embodiments of this system, an important part at the front of the Mux is upfront signal conditioning on Mux for improved signal-to-noise ratio. Embodiments may perform signal conditioning (such as range/gain control, integration, filtering, etc.) on vibration as well as other signal inputs up front before Mux switching to achieve the highest signal-to-noise ratio.

[0263] In example embodiments, in addition to providing a better signal, the multiplexer may provide a continuous monitor alarming feature. Truly continuous systems monitor every sensor all the time but tend to be expensive. Typical multiplexer systems only monitor a set number of channels at one time and switch from bank to bank of a larger set of sensors. As a result, the sensors not being currently collected are not being monitored; if a level increases the user may never know. In example embodiments, a multiplexer may have a continuous monitor alarming feature by placing circuitry on the multiplexer that can measure input channel levels against known alarm conditions even when the data acquisition ("DAQ") is not monitoring the input. In example embodiments, continuous monitoring Mux bypass offers a mechanism whereby channels not being currently sampled by the Mux system may be continuously monitored for significant alarm conditions via a number of trigger conditions using filtered peak-hold circuits or functionally

similar that are in turn passed on to the monitoring system in an expedient manner using hardware interrupts or other means. This, in essence, makes the system continuously monitor, although without the ability to instantly capture data on the problem like a true continuous system. In example embodiments, coupling this capability to alarm with adaptive scheduling techniques for continuous monitoring and the continuous monitoring system's software adapting and adjusting the data collection sequence based on statistics, analytics, data alarms and dynamic analysis may allow the system to quickly collect dynamic spectral data on the alarming sensor very soon after the alarm sounds.

[0264] Another restriction of typical multiplexers is that they may have a limited number of channels. In example embodiments, use of distributed complex programmable logic device ("CPLD") chips with dedicated bus for logic control of multiple Mux and data acquisition sections enables a CPLD to control multiple mux and DAQs so that there is no limit to the number of channels a system can handle. Interfacing to multiple types of predictive maintenance and vibration transducers requires a great deal of switching. This includes AC/DC coupling, 4-20 interfacing, integrated electronic piezoelectric transducer, channel power-down (for conserving op-amp power), single-ended or differential grounding options, and so on. Also required is the control of digital pots for range and gain control, switches for hardware integration, AA filtering and triggering. This logic can be performed by a series of CPLD chips strategically located for the tasks they control. A single giant CPLD requires long circuit routes with a great deal of density at the single giant CPLD. In example embodiments, distributed CPLDs not only address these concerns but offer a great deal of flexibility. A bus is created where each CPLD that has a fixed assignment has its own unique device address. In example embodiments, multiplexers and DAQs can stack together offering additional input and output channels to the system. For multiple boards (e.g., for multiple Mux boards), jumpers are provided for setting multiple addresses. In another example, three bits permit up to 8 boards that are jumper configurable. In example embodiments, a bus protocol is defined such that each CPLD on the bus can either be addressed individually or as a group.

[0265] Typical multiplexers may be limited to collecting only sensors in the same bank. For detailed analysis, this may be limiting as there is tremendous value in being able to simultaneously review data from sensors on the same machine. Current systems using conventional fixed bank multiplexers can only compare a limited number of channels (based on the number of channels per bank) that were assigned to a particular group at the time of installation. The only way to provide some flexibility is to either overlap channels or incorporate lots of redundancies in the system, both of which can add considerable expense (in some cases an exponential increase in cost versus flexibility). The simplest Mux design selects one of many inputs and routes it into a single output line. A banked design would consist of a group of these simple building blocks, each handling a fixed group of inputs and routing to its respective output. Typically, the inputs are not overlapping so that the input of one Mux grouping cannot be routed into another. Unlike conventional Mux chips which typically switch a fixed group or banks of a fixed selection of channels into a single output (e.g., in groups of 2, 4, 8, etc.), a cross point Mux allows the user to assign any input to any

output. Previously, crosspoint multiplexers were used for specialized purposes such as RGB digital video applications and were as a practical matter too noisy for analog applications such as vibration analysis; however more recent advances in the technology now make it feasible. Another advantage of the crosspoint Mux is the ability to disable outputs by putting them into a high impedance state. This is ideal for an output bus so that multiple Mux cards may be stacked, and their output buses joined together without the need for bus switches.

[0266] In example embodiments, this may be addressed by use of an analog crosspoint switch for collecting variable groups of vibration input channels and providing a matrix circuit so the system may access any set of eight channels from the total number of input sensors.

[0267] In example embodiments, the ability to control multiple multiplexers with use of distributed CPLD chips with dedicated bus for logic control of multiple Mux and data acquisition sections is enhanced with a hierarchical multiplexer which allows for multiple DAQ to collect data from multiple multiplexers. A hierarchical Mux may allow modular output of more channels, such as 16, 24 or more to multiple of eight channel card sets. In example embodiments, this allows for faster data collection as well as more channels of simultaneous data collection for more complex analysis. In example embodiments, the Mux may be configured slightly to make it portable and use data acquisition parking features, which turns SV3X DAQ into a protected system embodiment.

[0268] In example embodiments, once the signals leave the multiplexer and hierarchical Mux they move to the analog board where there are other enhancements. In example embodiments, power saving techniques may be used such as: power-down of analog channels when not in use; powering down of component boards; power-down of analog signal processing op-amps for non-selected channels; powering down channels on the mother and the daughter analog boards. The ability to power down component boards and other hardware by the low-level firmware for the DAQ system makes high-level application control with respect to power-saving capabilities relatively easy. Explicit control of the hardware is always possible but not required by default. In example embodiments, this power saving benefit may be of value to a protected system, especially if it is battery operated or solar powered.

[0269] In example embodiments, in order to maximize the signal to noise ratio and provide the best data, a peak-detector for auto-scaling routed into a separate A/D will provide the system the highest peak in each set of data so it can rapidly scale the data to that peak. For vibration analysis purposes, the built-in A/D converters in many microprocessors may be inadequate with regards to number of bits, number of channels or sampling frequency versus not slowing the microprocessor down significantly. Despite these limitations, it is useful to use them for the purposes of auto-scaling. In example embodiments, a separate A/D may be used that has reduced functionality and is cheaper. For each channel of input, after the signal is buffered (usually with the appropriate coupling: AC or DC) but before it is signal conditioned, the signal is fed directly into the microprocessor or low-cost A/D. Unlike the conditioned signal for which range, gain and filter switches are thrown, no switches are varied. This permits the simultaneous sampling of the auto-scaling data while the input data is signal conditioned, fed into a more robust external A/D, and

directed into on-board memory using direct memory access (DMA) methods where memory is accessed without requiring a CPU. This significantly simplifies the auto-scaling process by not having to throw switches and then allow for settling time, which greatly slows down the auto-scaling process. Furthermore, the data may be collected simultaneously, which assures the best signal-to-noise ratio. The reduced number of bits and other features is usually more than adequate for auto-scaling purposes. In example embodiments, improved integration using both analog and digital methods create an innovative hybrid integration which also improves or maintains the highest possible signal to noise ratio.

[0270] In example embodiments, a section of the analog board may allow routing of a trigger channel, either raw or buffered, into other analog channels. This may allow a user to route the trigger to any of the channels for analysis and trouble shooting. Systems may have trigger channels for the purposes of determining relative phase between various input data sets or for acquiring significant data without the needless repetition of unwanted input. In example embodiments, digitally controlled relays may be used to switch either the raw or buffered trigger signal into one of the input channels. It may be desirable to examine the quality of the triggering pulse because it may be corrupted for a variety of reasons including inadequate placement of the trigger sensor, wiring issues, faulty setup issues such as a dirty piece of reflective tape if using an optical sensor, and so on. The ability to look at either the raw or buffered signal may offer an excellent diagnostic or debugging vehicle. It also can offer some improved phase analysis capability by making use of the recorded data signal for various signal processing techniques such as variable speed filtering algorithms.

[0271] In example embodiments, once the signals leave the analog board, the signals move into the delta-sigma board where precise voltage reference for A/D zero reference offers more accurate direct current sensor data. The delta sigma's high speeds also provide for using higher input oversampling for delta-sigma A/D for lower sampling rate outputs to minimize antialiasing filter requirements. Lower oversampling rates can be used for higher sampling rates. For example, a 3rd order AA filter set for the lowest sampling requirement for 256 Hz (Fmax of 100 Hz) is then adequate for Fmax ranges of 200 and 500 Hz. Another higher-cutoff AA filter can then be used for Fmax ranges from 1 kHz and higher (with a secondary filter kicking in at 2.56x the highest sampling rate of 128 kHz). In example embodiments, a CPLD may be used as a clock-divider for a delta-sigma A/D to achieve lower sampling rates without the need for digital resampling. In example embodiments, a high-frequency crystal reference can be divided down to lower frequencies by employing a CPLD as a programmable clock divider. The accuracy of the divided down lower frequencies is even more accurate than the original source relative to their longer time periods. This also minimizes or removes the need for resampling processing by the delta-sigma A/D.

[0272] In example embodiments, the data then moves from the delta-sigma board to the Jennic™ board where phase relative to input and trigger channels using on-board timers may be digitally derived. In example embodiments, the Jennic™ board also has the ability to store calibration data and system maintenance repair history data in an on-board card set. In example embodiments, the

Jennic™ board will enable acquiring long blocks of data at high-sampling rate as opposed to multiple sets of data taken at different sampling rates so it can stream data and acquire long blocks of data for advanced analysis in the future.

[0273] In example embodiments, after the signal moves through the Jennic™ board it may then be transmitted to the computer. In example embodiments, the computer software will be used to add intelligence to the system starting with an expert system GUI. The GUI will offer a graphical expert system with simplified user interface for defining smart bands and diagnoses which facilitate anyone to develop complex analytics. In example embodiments, this user interface may revolve around smart bands, which are a simplified approach to complex yet flexible analytics for the general user. In example embodiments, the smart bands may pair with a self-learning neural network for an even more advanced analytical approach. In example embodiments, this system may use the machine's hierarchy for additional analytical insight. One critical part of predictive maintenance is the ability to learn from known information during repairs or inspections. In example embodiments, graphical approaches for back calculations may improve the smart bands and correlations based on a known fault or problem.

[0274] In example embodiments, there is a smart route which adapts which sensors it collects simultaneously in order to gain additional correlative intelligence. In example embodiments, smart operational data store ("ODS") allows the system to elect to gather data to perform operational deflection shape analysis in order to further examine the machinery condition. In example embodiments, adaptive scheduling techniques allow the system to change the scheduled data collected for full spectral analysis across a number (e.g., eight), of correlative channels. In example embodiments, the system may provide data to enable extended statistics capabilities for continuous monitoring as well as ambient local vibration for analysis that combines ambient temperature and local temperature and vibration levels changes for identifying machinery issues.

[0275] In example embodiments, a data acquisition device may be controlled by a personal computer (PC) to implement the desired data acquisition commands. In example embodiments, the DAQ box may be self-sufficient. and can acquire, process, analyze and monitor independent of external PC control. Embodiments may include secure digital (SD) card storage. In example embodiments, significant additional storage capability may be provided by utilizing an SD card. This may prove critical for monitoring applications where critical data may be stored permanently. Also, if a power failure should occur, the most recent data may be stored despite the fact that it was not off-loaded to another system.

[0276] A current trend has been to make DAQ systems as communicative as possible with the outside world usually in the form of networks including wireless. In the past it was common to use a dedicated bus to control a DAQ system with either a microprocessor or microcontroller/microprocessor paired with a PC. In example embodiments, a DAQ system may comprise one or more microprocessor/microcontrollers, specialized microcontrollers/microprocessors, or dedicated processors focused primarily on the communication aspects with the outside world. These include USB, Ethernet and wireless with the ability to provide an IP address or addresses in order to host a webpage. All communications with

the outside world are then accomplished using a simple text based menu. The usual array of commands (in practice more than a hundred) such as InitializeCard, AcquireData, StopAcquisition, RetrieveCalibration Info, and so on, would be provided.

[0277] In example embodiments, intense signal processing activities including resampling, weighting, filtering, and spectrum processing may be performed by dedicated processors such as field-programmable gate array ("FPGAs"), digital signal processor ("DSP"), microprocessors, micro-controllers, or a combination thereof. In example embodiments, this subsystem may communicate via a specialized hardware bus with the communication processing section. It will be facilitated with dual-port memory, semaphore logic, and so on. This embodiment will not only provide a marked improvement in efficiency but can significantly improve the processing capability, including the streaming of the data as well as other high-end analytical techniques. This negates the need for constantly interrupting the main processes which include the control of the signal conditioning circuits, triggering, raw data acquisition using the A/D, directing the A/D output to the appropriate on-board memory and processing that data.

[0278] Embodiments may include sensor overload identification. A need exists for monitoring systems to identify when the sensor is overloading. There may be situations involving high-frequency inputs that will saturate a standard 100 mv/g sensor (which is most commonly used in the industry) and having the ability to sense the overload improves data quality for better analysis. A monitoring system may identify when their system is overloading, but in example embodiments, the system may look at the voltage of the sensor to determine if the overload is from the sensor, enabling the user to get another sensor better suited to the situation, or gather the data again.

[0279] Embodiments may include radio frequency identification ("RFID") and an inclinometer or accelerometer on a sensor so the sensor can indicate what machine/bearing it is attached to and what direction such that the software can automatically store the data without the user input. In example embodiments, users could put the system on any machine or machines and the system would automatically set itself up and be ready for data collection in seconds.

[0280] Embodiments may include ultrasonic online monitoring by placing ultrasonic sensors inside transformers, motor control centers, breakers and the like and monitoring, via a sound spectrum, continuously looking for patterns that identify arcing, corona and other electrical issues indicating a break down or issue. Embodiments may include providing continuous ultrasonic monitoring of rotating elements and bearings of an energy production facility. In example embodiments, an analysis engine may be used in ultrasonic online monitoring as well as identifying other faults by combining the ultrasonic data with other parameters such as vibration, temperature, pressure, heat flux, magnetic fields, electrical fields, currents, voltage, capacitance, inductance, and combinations (e.g., simple ratios) of the same, among many others.

[0281] Embodiments of the methods and systems disclosed herein may include use of an analog crosspoint switch for collecting variable groups of vibration input channels. For vibration analysis, it is useful to obtain multiple channels simultaneously from vibration transducers mounted on different parts of a machine (or machines) in multiple directions. By obtaining the readings at the same time, for example, the relative phases of the inputs may be compared for the purpose of

diagnosing various mechanical faults. Other types of cross channel analyses such as cross-correlation, transfer functions, Operating Deflection Shape ("ODS") may also be performed.

[0282] Embodiments of the methods and systems disclosed herein may include precise voltage reference for A/D zero reference. Some A/D chips provide their own internal zero voltage reference to be used as a mid-scale value for external signal conditioning circuitry to ensure that both the A/D and external op-amps use the same reference. Although this sounds reasonable in principle, there are practical complications. In many cases these references are inherently based on a supply voltage using a resistor-divider. For many current systems, especially those whose power is derived from a PC via USB or similar bus, this provides for an unreliable reference, as the supply voltage will often vary quite significantly with load. This is especially true for delta-sigma A/D chips which necessitate increased signal processing. Although the offsets may drift together with load, a problem arises if one wants to calibrate the readings digitally. It is typical to modify the voltage offset expressed as counts coming from the A/D digitally to compensate for the DC drift. However, for this case, if the proper calibration offset is determined for one set of loading conditions, they will not apply for other conditions. An absolute DC offset expressed in counts will no longer be applicable. As a result, it becomes necessary to calibrate for all loading conditions which becomes complex, unreliable, and ultimately unmanageable. In example embodiments, an external voltage reference is used which is simply independent of the supply voltage to use as the zero offset.

[0283] In example embodiments, the system provides a phase-lock-loop band pass tracking filter method for obtaining slow-speed RPMs and phase for balancing purposes to remotely balance slow speed machinery, such as in paper mills, as well as offering additional analysis from its data. For balancing purposes, it is sometimes necessary to balance at very slow speeds. A typical tracking filter may be constructed based on a phase-lock loop or PLL design; however, stability and speed range are overriding concerns. In example embodiments, a number of digitally controlled switches are used for selecting the appropriate RC and damping constants. The switching can be done all automatically after measuring the frequency of the incoming tach signal. Embodiments of the methods and systems disclosed herein may include digital derivation of phase relative to input and trigger channels using on-board timers. In example embodiments, digital phase derivation uses digital timers to ascertain an exact delay from a trigger event to the precise start of data acquisition. This delay, or offset, then, is further refined using interpolation methods to obtain an even more precise offset which is then applied to the analytically determined phase of the acquired data such that the phase is "in essence" an absolute phase with precise mechanical meaning useful for among other things, one-shot balancing, alignment analysis, and so on.

[0284] Embodiments of the methods and systems disclosed herein may include signal processing firmware/hardware. In example embodiments, long blocks of data may be acquired at high-sampling rate as opposed to multiple sets of data taken at different sampling rates. Typically, in modern route collection for vibration analysis, it is customary to collect data at a fixed sampling rate with a specified data length. The sampling rate and data length may vary from route point to point based on the specific mechanical analysis requirements at hand. For example, a motor may require a relatively low sampling rate with high resolution to distinguish running speed harmonics

from line frequency harmonics. The practical trade-off here though is that it takes more collection time to achieve this improved resolution. In contrast, some high-speed compressors or gear sets require much higher sampling rates to measure the amplitudes of relatively higher frequency data although the precise resolution may not be as necessary. Ideally, however, it would be better to collect a very long sample length of data at a very high-sampling rate. When digital acquisition devices were first popularized in the early 1980s, the A/D sampling, digital storage, and computational abilities were not close to what they are today, so compromises were made between the time required for data collection and the desired resolution and accuracy. It was because of this limitation that some analysts in the field even refused to give up their analog tape recording systems, which did not suffer as much from these same digitizing drawbacks. A few hybrid systems were employed that would digitize the play back of the recorded analog data at multiple sampling rates and lengths desired, though these systems were admittedly less automated. The more common approach, as mentioned earlier, is to balance data collection time with analysis capability and digitally acquire the data blocks at multiple sampling rates and sampling lengths and digitally store these blocks separately. In example embodiments, a long data length of data can be collected at the highest practical sampling rate (e.g., 102.4 kHz; corresponding to a 40 kHz Fmax) and stored. This long block of data can be acquired in the same amount of time as the shorter length of the lower sampling rates utilized by a priori methods so that there is no effective delay added to the sampling at the measurement point, always a concern in route collection. In example embodiments, analog tape recording of data is digitally simulated with such a precision that it can be in effect considered continuous or "analog" for many purposes, including for purposes of embodiments of the present disclosure, except where context indicates otherwise.

[0285] Embodiments of the methods and systems disclosed herein may include storage of calibration data and maintenance history on-board card sets. Many data acquisition devices which rely on interfacing to a PC to function store their calibration coefficients on the PC. This is especially true for complex data acquisition devices whose signal paths are many and therefore whose calibration tables can be quite large. In example embodiments, calibration coefficients are stored in flash memory which will remember this data or any other significant information for that matter, for all practical purposes, permanently. This information may include nameplate information such as serial numbers of individual components, firmware or software version numbers, maintenance history, and the calibration tables. In example embodiments, no matter which computer the box is ultimately connected to, the DAQ box remains calibrated and continues to hold all of this critical information. The PC or external device may poll for this information at any time for implantation or information exchange purposes.

[0286] Embodiments of the methods and systems disclosed herein may include rapid route creation taking advantage of hierarchical templates. In the field of vibration monitoring, as well as parametric monitoring in general, it is necessary to establish in a database or functional equivalent the existence of data monitoring points. These points are associated with a variety of attributes including the following categories: transducer attributes, data collection settings, machinery parameters and operating parameters. The transducer attributes would include probe type, probe

mounting type and probe mounting direction or axis orientation. Data collection attributes associated with the measurement would involve a sampling rate, data length, integrated electronic piezoelectric probe power and coupling requirements, hardware integration requirements, 4-20 or voltage interfacing, range and gain settings (if applicable), filter requirements, and so on. Machinery parametric requirements relative to the specific point would include such items as operating speed, bearing type, bearing parametric data which for a rolling element bearing includes the pitch diameter, number of balls, inner race, and outer-race diameters. For a tilting pad bearing, this would include the number of pads and so on. For measurement points on a piece of equipment such as a gearbox, needed parameters would include, for example, the number of gear teeth on each of the gears. For induction motors, it would include the number of rotor bars and poles; for compressors, the number of blades and/or vanes; for fans, the number of blades. For belt/pulley systems, the number of belts as well as the relevant belt-passing frequencies may be calculated from the dimensions of the pulleys and pulley center-to-center distance. For measurements near couplings, the coupling type and number of teeth in a geared coupling may be necessary, and so on. Operating parametric data would include operating load, which may be expressed in megawatts, flow (either air or fluid), percentage, horsepower, feet-per-minute, and so on. Operating temperatures both ambient and operational, pressures, humidity, and so on, may also be relevant. As can be seen, the setup information required for an individual measurement point can be quite large. It is also crucial to performing any legitimate analysis of the data. Machinery, equipment, and bearing specific information are essential for identifying fault frequencies as well as anticipating the various kinds of specific faults to be expected. The transducer attributes as well as data collection parameters are vital for properly interpreting the data along with providing limits for the type of analytical techniques suitable. The traditional means of entering this data has been manual and quite tedious, usually at the lowest hierarchical level (for example, at the bearing level with regards to machinery parameters), and at the transducer level for data collection setup information. It cannot be stressed enough, however, the importance of the hierarchical relationships necessary to organize data— both for analytical and interpretive purposes as well as the storage and movement of data. Here, we are focusing primarily on the storage and movement of data. By its nature, the aforementioned setup information is extremely redundant at the level of the lowest hierarchies; however, because of its strong hierarchical nature, it can be stored quite efficiently in that form. In example embodiments, hierarchical nature can be utilized when copying data in the form of templates. As an example, hierarchical storage structure suitable for many purposes is defined from general to specific of company, plant or site, unit or process, machine, equipment, shaft element, bearing, and transducer. It is much easier to copy data associated with a particular machine, piece of equipment, shaft element or bearing than it is to copy only at the lowest transducer level. In example embodiments, the system not only stores data in this hierarchical fashion, but robustly supports the rapid copying of data using these hierarchical templates. Similarity of elements at specific hierarchical levels lends itself to effective data storage in hierarchical format. For example, so many machines have common elements such as motors, gearboxes, compressors, belts, fans, and so on. More specifically, many motors can be easily

classified as induction, DC, fixed or variable speed. Many gearboxes can be grouped into commonly occurring groupings such as input/output, input pinion/intermediate pinion/output pinion, 4-posters, and so on. Within a plant or company, there are many similar types of equipment purchased and standardized on for both cost and maintenance reasons. This results in an enormous overlapping of similar types of equipment and, as a result, offers a great opportunity for taking advantage of a hierarchical template approach.

[0287] Embodiments of the methods and systems disclosed herein may include smart bands. Smart bands refer to any processed signal characteristics derived from any dynamic input or group of inputs for the purposes of analyzing the data and achieving the correct diagnoses. Furthermore, smart bands may even include mini or relatively simple diagnoses for the purposes of achieving a more robust and complex one. Historically, in the field of mechanical vibration analysis, Alarm Bands have been used to define spectral frequency bands of interest for the purposes of analyzing and/or trending significant vibration patterns. The Alarm Band typically consists of a spectral (amplitude plotted against frequency) region defined between a low and high frequency border. The amplitude between these borders is summed in the same manner for which an overall amplitude is calculated. A Smart Band is more flexible in that it not only refers to a specific frequency band but can also refer to a group of spectral peaks such as the harmonics of a single peak, a true-peak level or crest factor derived from a time waveform, an overall derived from a vibration envelope spectrum or other specialized signal analysis technique or a logical combination (AND, OR, XOR, etc.) of these signal attributes. In addition, a myriad assortment of other parametric data, including system load, motor voltage and phase information, bearing temperature, flow rates, and the like, can likewise be used as the basis for forming additional smart bands. In example embodiments, Smart Band symptoms may be used as building blocks for an expert system whose engine would utilize these inputs to derive diagnoses. Some of these mini-diagnoses may then in turn be used as Smart-Band symptoms (smart bands can include even diagnoses) for more generalized diagnoses.

[0288] Embodiments of the methods and systems disclosed herein may include a neural net expert system using smart bands. Typical vibration analysis engines are rule-based (i.e., they use a list of expert rules which, when met, trigger specific diagnoses). In contrast, a neural approach utilizes the weighted triggering of multiple input stimuli into smaller analytical engines or neurons which in turn feed a simplified weighted output to other neurons. The output of these neurons can be also classified as smart bands which in turn feed other neurons. This produces a more layered approach to expert diagnosing as opposed to the one-shot approach of a rule-based system. In example embodiments, the expert system utilizes this neural approach using smart bands; however, it does not preclude rule-based diagnoses being reclassified as smart bands as further stimuli to be utilized by the expert system. From this point-of-view, it can be overviewed as a hybrid approach, although at the highest level it is essentially neural.

[0289] Embodiments of the methods and systems disclosed herein may include use of database hierarchy in analysis smart band symptoms and diagnoses may be assigned to various hierarchical database levels. For example, a smart band may be called "Looseness" at the bearing level, trigger

"Looseness" at the equipment level, and trigger "Looseness" at the machine level. Another example would be having a smart band diagnosis called "Horizontal Plane Phase Flip" across a coupling and generate a smart band diagnosis of "Vertical Coupling Misalignment" at the machine level.

[0290] Embodiments of the methods and systems disclosed herein may include expert system GUIs. In example embodiments, the system undertakes a graphical approach to defining smart bands and diagnoses for the expert system. The entry of symptoms, rules, or more generally smart bands for creating a particular machine diagnosis, may be tedious and time consuming. One means of making the process more expedient and efficient is to provide a graphical means by use of wiring. The proposed graphical interface consists of four major components: a symptom parts bin, diagnoses bin, tools bin, and graphical wiring area ("GWA"). In example embodiments, a symptom parts bin includes various spectral, waveform, envelope and any type of signal processing characteristic or grouping of characteristics such as a spectral peak, spectral harmonic, waveform true-peak, waveform crest-factor, spectral alarm band, and so on. Each part may be assigned additional properties. For example, a spectral peak part may be assigned a frequency or order (multiple) of running speed. Some parts may be pre-defined or user defined such as a 1x, 2x, 3x running speed, 1x, 2x, 3x gear mesh, 1x, 2x, 3x blade pass, number of motor rotor bars x running speed, and so on.

[0291] In example embodiments, the diagnoses bin includes various pre-defined as well as user-defined diagnoses such as misalignment, imbalance, looseness, bearing faults, and so on. Like parts, diagnoses may also be used as parts for the purposes of building more complex diagnoses. In example embodiments, the tools bin includes logical operations such as AND, OR, XOR, etc. or other ways of combining the various parts listed above such as Find Max, Find Min, Interpolate, Average, other Statistical Operations, etc. In example embodiments, a graphical wiring area includes parts from the parts bin or diagnoses from the diagnoses bin and may be combined using tools to create diagnoses. The various parts, tools and diagnoses will be represented with icons which are simply graphically wired together in the desired manner.

[0292] Embodiments of the methods and systems disclosed herein may include a graphical approach for back-calculation definition. In example embodiments, the expert system also provides the opportunity for the system to learn. If one already knows that a unique set of stimuli or smart bands corresponds to a specific fault or diagnosis, then it is possible to back-calculate a set of coefficients that when applied to a future set of similar stimuli would arrive at the same diagnosis. In example embodiments, if there are multiple sets of data, a best-fit approach may be used. Unlike the smart band GUI, this embodiment will self-generate a wiring diagram. In example embodiments, the user may tailor the back-propagation approach settings and use a database browser to match specific sets of data with the desired diagnoses. In example embodiments, the desired diagnoses may be created or custom tailored with a smart band GUI. In example embodiments, after that, a user may press the GENERATE button and a dynamic wiring of the symptom-to-diagnosis may appear on the screen as it works through the algorithms to achieve the best fit. In example embodiments, when complete, a variety of statistics are presented which detail

how well the mapping process proceeded. In some cases, no mapping may be achieved if, for example, the input data was all zero or the wrong data (mistakenly assigned) and so on. Embodiments of the methods and systems disclosed herein may include bearing analysis methods. In example embodiments, bearing analysis methods may be used in conjunction with a computer aided design ("CAD"), predictive deconvolution, minimum variance distortionless response ("MVDR") and spectrum sum-of-harmonics.

[0293] In recent years, there has been a strong drive to save power which has resulted in an influx of variable frequency drives and variable speed machinery. In example embodiments, a bearing analysis method is provided. In example embodiments, torsional vibration detection and analysis is provided utilizing transitory signal analysis to provide an advanced torsional vibration analysis for a more comprehensive way to diagnose machinery where torsional forces are relevant (such as machinery with rotating components). Due primarily to the decrease in cost of motor speed control systems, as well as the increased cost and consciousness of energy-usage, it has become more economically justifiable to take advantage of the potentially vast energy savings of load control. Unfortunately, one frequently overlooked design aspect of this issue is that of vibration. When a machine is designed to run at only one speed, it is far easier to design the physical structure accordingly so as to avoid mechanical resonances both structural and torsional, each of which can dramatically shorten the mechanical health of a machine. This would include such structural characteristics as the types of materials to use, their weight, stiffening member requirements and placement, bearing types, bearing location, base support constraints, etc. Even with machines running at one speed, designing a structure so as to minimize vibration can prove a daunting task, potentially requiring computer modeling, finite-element analysis, and field testing. By throwing variable speeds into the mix, in many cases, it becomes impossible to design for all desirable speeds. The problem then becomes one of minimization, e.g., by speed avoidance. This is why many modern motor controllers are typically programmed to skip or quickly pass through specific speed ranges or bands. Embodiments may include identifying speed ranges in a vibration monitoring system. Non-torsional, structural resonances are typically fairly easy to detect using conventional vibration analysis techniques. However, this is not the case for torsion. One special area of current interest is the increased incidence of torsional resonance problems, apparently due to the increased torsional stresses of speed change as well as the operation of equipment at torsional resonance speeds. Unlike non-torsional structural resonances which generally manifest their effect with dramatically increased casing or external vibration, torsional resonances generally show no such effect. In the case of a shaft torsional resonance, the twisting motion induced by the resonance may only be discernible by looking for speed and/or phase changes. The current standard methodology for analyzing torsional vibration involves the use of specialized instrumentation. Methods and systems disclosed herein allow analysis of torsional vibration without such specialized instrumentation. This may consist of shutting the machine down and employing the use of strain gauges and/or other special fixturing such as speed encoder plates and/or gears. Friction wheels are another alternative, but they typically require manual implementation and a specialized analyst. In general, these techniques can be prohibitively expensive and/or inconvenient. An

increasing prevalence of continuous vibration monitoring systems due to decreasing costs and increasing convenience (e.g., remote access) exists. In example embodiments, there is an ability to discern torsional speed and/or phase variations with just the vibration signal. In example embodiments, transient analysis techniques may be utilized to distinguish torsionally induced vibrations from mere speed changes due to process control. In example embodiments, factors for discernment might focus on one or more of the following aspects: the rate of speed change due to variable speed motor control would be relatively slow, sustained and deliberate; torsional speed changes would tend to be short, impulsive and not sustained; torsional speed changes would tend to be oscillatory, most likely decaying exponentially, process speed changes would not; and smaller speed changes associated with torsion relative to the shaft's rotational speed which suggest that monitoring phase behavior would show the quick or transient speed bursts in contrast to the slow phase changes historically associated with ramping a machine's speed up or down (as typified with Bode or Nyquist plots).

[0294] Embodiments of the methods and systems disclosed herein may include improved integration using both analog and digital methods. When a signal is digitally integrated using software, essentially the spectral low-end frequency data has its amplitude multiplied by a function which quickly blows up as it approaches zero and creates what is known in the industry as a "ski-slope" effect. The amplitude of the ski-slope is essentially the noise floor of the instrument. The simple remedy for this is the traditional hardware integrator, which can perform at signal-to-noise ratios much greater than that of an already digitized signal. It can also limit the amplification factor to a reasonable level so that multiplication by very large numbers is essentially prohibited. However, at high frequencies where the frequency becomes large, the original amplitude which may be well above the noise floor is multiplied by a very small number (1/f) that plunges it well below the noise floor. The hardware integrator has a fixed noise floor that although low floor does not scale down with the now lower amplitude high-frequency data. In contrast, the same digital multiplication of a digitized high-frequency signal also scales down the noise floor proportionally. In example embodiments, hardware integration may be used below the point of unity gain where (at a value usually determined by units and/or desired signal to noise ratio based on gain) and software integration may be used above the value of unity gain to produce an ideal result. In example embodiments, this integration is performed in the frequency domain. In example embodiments, the resulting hybrid data can then be transformed back into a waveform which should be far superior in signal-to-noise ratio when compared to either hardware integrated or software integrated data. In example embodiments, the strengths of hardware integration are used in conjunction with those of digital software integration to achieve the maximum signal-to-noise ratio. In example embodiments, the first order gradual hardware integrator high pass filter along with curve fitting allow some relatively low frequency data to get through while reducing or eliminating the noise, allowing very useful analytical data that steep filters kill to be salvaged.

[0295] Embodiments of the methods and systems disclosed herein may include adaptive scheduling techniques for continuous monitoring. Continuous monitoring is often performed with an up-front Mux whose purpose it is to select a few channels of data among many to feed the

hardware signal processing, A/D, and processing components of a DAQ system. This is done primarily out of practical cost considerations. The tradeoff is that all of the points are not monitored continuously (although they may be monitored to a lesser extent via alternative hardware methods). In example embodiments, multiple scheduling levels are provided. In example embodiments, at the lowest level, which is continuous for the most part, all of the measurement points will be cycled through in round-robin fashion. For example, if it takes 30 seconds to acquire and process a measurement point and there are 30 points, then each point is serviced once every 15 minutes; however, if a point should alarm by whatever criteria the user selects, its priority level can be increased so that it is serviced more often. As there can be multiple grades of severity for each alarm, so can there me multiple levels of priority with regards to monitoring. In example embodiments, more severe alarms will be monitored more frequently. In example embodiments, a number of additional high-level signal processing techniques can be applied at less frequent intervals. Embodiments may take advantage of the increased processing power of a PC and the PC can temporarily suspend the round-robin route collection (with its multiple tiers of collection) process and stream the required amount of data for a point of its choosing. Embodiments may include various advanced processing techniques such as envelope processing, wavelet analysis, as well as many other signal processing techniques. In example embodiments, after acquisition of this data, the DAQ card set will continue with its route at the point it was interrupted. In example embodiments, various PC scheduled data acquisitions will follow their own schedules which will be less frequency than the DAQ card route. They may be set up hourly, daily, by number of route cycles (for example, once every 10 cycles) and also increased scheduling-wise based on their alarm severity priority or type of measurement (e.g., motors may be monitored differently than fans).

[0296] Embodiments of the methods and systems disclosed herein may include data acquisition parking features. In example embodiments, a data acquisition box used for route collection, real time analysis and in general as an acquisition instrument can be detached from its PC (tablet or otherwise) and powered by an external power supply or suitable battery. In example embodiments, the data collector still retains continuous monitoring capability and its on-board firmware can implement dedicated monitoring functions for an extended period of time or can be controlled remotely for further analysis. Embodiments of the methods and systems disclosed herein may include extended statistical capabilities for continuous monitoring.

[0297] Embodiments of the methods and systems disclosed herein may include ambient sensing plus local sensing plus vibration for analysis. In example embodiments, ambient environmental temperature and pressure, sensed temperature and pressure may be combined with long/medium term vibration analysis for prediction of any of a range of conditions or characteristics. Variants may add infrared sensing, infrared thermography, ultrasound, and many other types of sensors and input types in combination with vibration or with each other. Embodiments of the methods and systems disclosed herein may include a smart route. In example embodiments, the continuous monitoring system's software will adapt/adjust the data collection sequence based on statistics, analytics, data alarms and dynamic analysis. Typically, the route is set based on the channels the sensors are attached to. In example embodiments, with the crosspoint switch, the Mux can combine

any input Mux channels to the (e.g., eight) output channels. In example embodiments, as channels go into alarm or the system identifies key deviations, it will pause the normal route set in the software to gather specific simultaneous data, from the channels sharing key statistical changes, for more advanced analysis. Embodiments include conducting a smart ODS or smart transfer function.

[0298] Embodiments of the methods and systems disclosed herein may include smart ODS and one or more transfer functions. In example embodiments, due to a system's multiplexer and crosspoint switch, an ODS, a transfer function, or other special tests on all the vibration sensors attached to a machine/structure can be performed and show exactly how the machine's points are moving in relationship to each other. In example embodiments, 40-50 kHz and longer data lengths (e.g., at least one minute) may be streamed, which may reveal different information than what a normal ODS or transfer function will show. In example embodiments, the system will be able to determine, based on the data/statistics/analytics to use, the smart route feature that breaks from the standard route and conducts an ODS across a machine, structure or multiple machines and structures that might show a correlation because the conditions/data directs it. In example embodiments, for the transfer functions there may be an impact hammer used on one channel and then compared against other vibration sensors on the machine. In example embodiments, the system may use the condition changes such as load, speed, temperature or other changes in the machine or system to conduct the transfer function. In example embodiments, different transfer functions may be compared to each other over time. In example embodiments, difference transfer functions may be strung together like a movie that may show how the machinery fault changes, such as a bearing that could show how it moves through the four stages of bearing failure and so on. Embodiments of the methods and systems disclosed herein may include a hierarchical Mux.

[0299] With reference to Figure 8, the present disclosure generally includes digitally collecting or streaming waveform data 2010 from a machine 2020 whose operational speed can vary from relatively slow rotational or oscillational speeds to much higher speeds in different situations. The waveform data 2010, at least on one machine, may include data from a single axis sensor 2030 mounted at an unchanging reference location 2040 and from a three-axis sensor 2050 mounted at changing locations (or located at multiple locations), including location 2052. In example embodiments, the waveform data 2010 can be vibration data obtained simultaneously from each sensor 2030, 2050 in a gap-free format for a duration of multiple minutes with maximum resolvable frequencies sufficiently large to capture periodic and transient impact events. By way of this example, the waveform data 2010 can include vibration data that can be used to create an operational deflecting shape. It can also be used, as needed, to diagnose vibrations from which a machine repair solution can be prescribed.

[0300] In example embodiments, the machine 2020 can further include a housing 2100 that can contain a drive motor 2110 that can drive a shaft 2120. The shaft 2120 can be supported for rotation or oscillation by a set of bearings 2130, such as including a first bearing 2140 and a second bearing 2150. A data collection module 2160 can connect to (or be resident on) the machine 2020. In one example, the data collection module 2160 can be located and accessible through a cloud network

facility 2170, can collect the waveform data 2010 from the machine 2020, and deliver the waveform data 2010 to a remote location. A working end 2180 of the drive shaft 2120 of the machine 2020 can drive a windmill, a fan, a pump, a drill, a gear system, a drive system, or other working element, as the techniques described herein can apply to a wide range of machines, equipment, tools, or the like that include rotating or oscillating elements. In other instances, a generator can be substituted for the motor 2110, and the working end of the drive shaft 2120 can direct rotational energy to the generator to generate power, rather than consume it.

[0301] In example embodiments, the waveform data 2010 can be obtained using a predetermined route format based on the layout of the machine 2020. The waveform data 2010 may include data from the single axis sensor 2030 and the three-axis sensor 2050. The single-axis sensor 2030 can serve as a reference probe with its one channel of data and can be fixed at the unchanging location 2040 on the machine under survey. The three-axis sensor 2050 can serve as a tri-axial probe (e.g., three orthogonal axes) with its three channels of data and can be moved along a predetermined diagnostic route format from one test point to the next test point. In one example, both sensors 2030, 2050 can be mounted manually to the machine 2020 and can connect to a separate portable computer in certain service examples. The reference probe can remain at one location while the user can move the tri-axial vibration probe along the predetermined route, such as from bearing-to-bearing on a machine. In this example, the user is instructed to locate the sensors at the predetermined locations to complete the survey (or portion thereof) of the machine.

[0302] With reference to Figure 9, a portion of an exemplary machine 2200 is shown having a tri-axial sensor 2210 mounted to a location 2220 associated with a motor bearing of the machine 2200 with an output shaft 2230 and output member 2240 in accordance with the present disclosure.

[0303] In further examples, the sensors and data acquisition modules and equipment can be integral to, or resident on, the rotating machine. By way of these examples, the machine can contain many single axis sensors and many tri-axial sensors at predetermined locations. The sensors can be originally installed equipment and provided by the original equipment manufacturer or installed at a different time in a retrofit application. The data collection module 2160, or the like, can select and use one single axis sensor and obtain data from it exclusively during the collection of waveform data 2010 while moving to each of the tri-axial sensors. The data collection module 2160 can be resident on the machine 2020 and/or connect via the cloud network facility 2170.

[0304] With reference to Figure 8, the various embodiments include collecting the waveform data 2010 by digitally recording locally, or streaming over, the cloud network facility 2170. The waveform data 2010 can be collected so as to be gap-free with no interruptions and, in some respects, can be similar to an analog recording of waveform data. The waveform data 2010 from all of the channels can be collected for one to two minutes depending on the rotating or oscillating speed of the machine being monitored. In example embodiments, the data sampling rate can be at a relatively high-sampling rate relative to the operating frequency of the machine 2020.

[0305] In example embodiments, a second reference sensor can be used, and a fifth channel of data can be collected. As such, the single-axis sensor can be the first channel and tri-axial vibration can occupy the second, the third, and the fourth data channels. This second reference sensor, like the

first, can be a single axis sensor, such as an accelerometer. In example embodiments, the second reference sensor, like the first reference sensor, can remain in the same location on the machine for the entire vibration survey on that machine. The location of the first reference sensor (i.e., the single axis sensor) may be different than the location of the second reference sensors (i.e., another single axis sensor). In certain examples, the second reference sensor can be used when the machine has two shafts with different operating speeds, with the two reference sensors being located on the two different shafts. In accordance with this example, further single-axis reference sensors can be employed at additional but different unchanging locations associated with the rotating machine.

[0306] In example embodiments, the waveform data can be transmitted electronically in a gap-free free format at a significantly high rate of sampling for a relatively longer period of time. In one example, the period of time is 60 seconds to 120 seconds. In another example, the rate of sampling is 100 kHz with a maximum resolvable frequency (Fmax) of 40 kHz. It will be appreciated in light of this disclosure that the waveform data can be shown to approximate more closely some of the wealth of data available from previous instances of analog recording of waveform data.

[0307] In example embodiments, sampling, band selection, and filtering techniques can permit one or more portions of a long stream of data (i.e., one to two minutes in duration) to be under sampled or over sampled to realize varying effective sampling rates. To this end, interpolation and decimation can be used to further realize varying effective sampling rates. For example, oversampling may be applied to frequency bands that are proximal to rotational or oscillational operating speeds of the sampled machine, or to harmonics thereof, as vibration effects may tend to be more pronounced at those frequencies across the operating range of the machine. In example embodiments, the digitally-sampled data set can be decimated to produce a lower sampling rate. It will be appreciated in light of the disclosure that decimate in this context can be the opposite of interpolate. In example embodiments, decimating the data set can include first applying a low-pass filter to the digitally-sampled data set and then undersampling the data set.

[0308] In one example, a sample waveform at 100 Hz can be undersampled at every tenth point of the digital waveform to produce an effective sampling rate of 10 Hz, but the remaining nine points of that portion of the waveform are effectively discarded and not included in the modeling of the sample waveform. Moreover, this type of bare undersampling can create ghost frequencies due to the undersampling rate (i.e., 10 Hz) relative to the 100 Hz sample waveform.

[0309] Most hardware for analog-to-digital conversions uses a sample-and-hold circuit that can charge up a capacitor for a given amount of time such that an average value of the waveform is determined over a specific change in time. It will be appreciated in light of the disclosure that the value of the waveform over the specific change in time is not linear but more similar to a cardinal sinusoidal ("sinc") function; therefore, it can be shown that more emphasis can be placed on the waveform data at the center of the sampling interval with exponential decay of the cardinal sinusoidal signal occurring from its center.

[0310] By way of the above example, the sample waveform at 100 Hz can be hardware-sampled at 10 Hz and therefore each sampling point is averaged over 100 milliseconds (e.g., a signal sampled at 100 Hz can have each point averaged over 10 milliseconds). In contrast to the effective

discarding of nine out of the ten data points of the sampled waveform as discussed above, the present disclosure can include weighing adjacent data. The adjacent data can refer to the sample points that were previously discarded and the one remaining point that was retained. In one example, a low pass filter can average the adjacent sample data linearly, i.e., determining the sum of every ten points and then dividing that sum by ten. In a further example, the adjacent data can be weighted with a sinc function. The process of weighting the original waveform with the sinc function can be referred to as an impulse function, or can be referred to in the time domain as a convolution.

[0311] The present disclosure can be applicable to not only digitizing a waveform signal based on a detected voltage, but can also be applicable to digitizing waveform signals based on current waveforms, vibration waveforms, and image processing signals including video signal rasterization. In one example, the resizing of a window on a computer screen can be decimated, albeit in at least two directions. In these further examples, it will be appreciated that undersampling by itself can be shown to be insufficient. To that end, oversampling or upsampling by itself can similarly be shown to be insufficient, such that interpolation can be used like decimation but in lieu of only undersampling by itself.

[0312] It will be appreciated in light of the disclosure that interpolation in this context can refer to first applying a low pass filter to the digitally-sampled waveform data and then upsampling the waveform data. It will be appreciated in light of the disclosure that real-world examples can often require the use of use non-integer factors for decimation or interpolation, or both. To that end, the present disclosure includes interpolating and decimating sequentially in order to realize a non-integer factor rate for interpolating and decimating. In one example, interpolating and decimating sequentially can define applying a low-pass filter to the sample waveform, then interpolating the waveform after the low-pass filter, and then decimating the waveform after the interpolation. In example embodiments, the vibration data can be looped to purposely emulate conventional tape recorder loops, with digital filtering techniques used with the effective splice to facilitate longer analyses. It will be appreciated in light of the disclosure that the above techniques do not preclude waveform, spectrum, and other types of analyses from being processed and displayed with a GUI of the user at the time of collection. It will be appreciated in light of the disclosure that newer systems can permit this functionality to be performed in parallel to the high-performance collection of the raw waveform data.

[0313] With respect to time of collection issues, it will be appreciated that older systems using the compromised approach of improving data resolution, by collecting at different sampling rates and data lengths, do not in fact save as much time as expected. To that end, every time the data acquisition hardware is stopped and started, latency issues can be created, especially when there is hardware auto-scaling performed. The same can be true with respect to data retrieval of the route information (i.e., test locations) that is often in a database format and can be exceedingly slow. The storage of the raw data in bursts to disk (whether solid state or otherwise) can also be undesirably slow.

[0314] In contrast, the many embodiments include digitally streaming the waveform data 2010, as disclosed herein, and also enjoying the benefit of needing to load the route parameter information while setting the data acquisition hardware only once. Because the waveform data 2010 is streamed to only one file, there is no need to open and close files, or switch between loading and writing operations with the storage medium. It can be shown that the collection and storage of the waveform data 2010, as described herein, can be shown to produce relatively more meaningful data in significantly less time than the traditional batch data acquisition approach. An example of this includes an electric motor about which waveform data can be collected with a data length of 4K points (i.e., 4,096) for sufficiently high resolution in order to, among other things, distinguish electrical sideband frequencies. For fans or blowers, a reduced resolution of 1K (i.e., 1,024) can be used. In certain instances, 1K can be the minimum waveform data length requirement. The sampling rate can be 1,280 Hz and that equates to an Fmax of 500 Hz. It will be appreciated in light of the disclosure that oversampling by an industry standard factor of 2.56 can satisfy the necessary two-times (2x) oversampling for the Nyquist Criterion with some additional leeway that can accommodate anti-aliasing filter-rolloff. The time to acquire this waveform data would be 1,024 points at 1,280 hertz, which are 800 milliseconds.

[0315] To improve accuracy, the waveform data can be averaged. Eight averages can be used with, for example, fifty percent overlap. This would extend the time from 800 milliseconds to 3.6 seconds, which is equal to 800 msec x 8 averages x 0.5 (overlap ratio) + 0.5 x 800 msec (non-overlapped head and tail ends). After collection at Fmax = 500 Hz waveform data, a higher sampling rate can be used. In one example, ten times (10x) the previous sampling rate can be used and Fmax = 10 kHz. By way of this example, eight averages can be used with fifty percent (50%) overlap to collect waveform data at this higher rate that can amount to a collection time of 360 msec or 0.36 seconds. It will be appreciated in light of the disclosure that it can be necessary to read the hardware collection parameters for the higher sampling rate from the route list, as well as permit hardware auto-scaling, or the resetting of other necessary hardware collection parameters, or both. To that end, a few seconds of latency can be added to accommodate the changes in sampling rate. In other instances, introducing latency can accommodate hardware autoscaling and changes to hardware collection parameters that can be required when using the lower sampling rate disclosed herein. In addition to accommodating the change in sampling rate, additional time is needed for reading the route point information from the database (i.e., where to monitor and where to monitor next), displaying the route information, and processing the waveform data. Moreover, display of the waveform data and/or associated spectra can also consume significant time. In light of the above, 15 seconds to 20 seconds can elapse while obtaining waveform data at each measurement point.

[0316] In further examples, additional sampling rates can be added but this can make the total amount time for the vibration survey even longer because time adds up from changeover time from one sampling rate to another and from the time to obtain additional data at different sampling rate. In one example, a lower sampling rate is used, such as a sampling rate of 128 Hz where Fmax = 50 Hz. By way of this example, the vibration survey would, therefore, require an additional 36

seconds for the first set of averaged data at this sampling rate, in addition to others mentioned above, and consequently the total time spent at each measurement point increases even more dramatically. Further embodiments include using similar digital streaming of gap free waveform data as disclosed herein for use with wind turbines and other machines that can have relatively slow speed rotating or oscillating systems. In many examples, the waveform data collected can include long samples of data at a relatively high-sampling rate. In one example, the sampling rate can be 100 kHz and the sampling duration can be for two minutes on all of the channels being recorded. In many examples, one channel can be for the single axis reference sensor and three more data channels can be for the tri-axial three channel sensor. It will be appreciated in light of the disclosure that the long data length can be shown to facilitate detection of extremely low frequency phenomena. The long data length can also be shown to accommodate the inherent speed variability in wind turbine operations. Additionally, the long data length can further be shown to provide the opportunity for using numerous averages such as those discussed herein, to achieve very high spectral resolution, and to make feasible tape loops for certain spectral analyses. Many multiple advanced analytical techniques can now become available because such techniques can use the available long uninterrupted length of waveform data in accordance with the present disclosure.

[0317] It will also be appreciated in light of the disclosure that the simultaneous collection of waveform data from multiple channels can facilitate performing transfer functions between multiple channels. Moreover, the simultaneous collection of waveform data from multiple channels facilitates establishing phase relationships across the machine so that more sophisticated correlations can be utilized by relying on the fact that the waveforms from each of the channels are collected simultaneously. In other examples, more channels in the data collection can be used to reduce the time it takes to complete the overall vibration survey by allowing for simultaneous acquisition of waveform data from multiple sensors that otherwise would have to be acquired, in a subsequent fashion, moving sensor to sensor in the vibration survey.

[0318] The present disclosure includes the use of at least one of the single-axis reference probes on one of the channels to allow for acquisition of relative phase comparisons between channels. The reference probe can be an accelerometer or other type of transducer that is not moved and, therefore, fixed at an unchanging location during the vibration survey of one machine. Multiple reference probes can each be deployed at suitable locations fixed in place (i.e., at unchanging locations) throughout the acquisition of vibration data during the vibration survey. In certain examples, up to seven reference probes can be deployed depending on the capacity of the data collection module 2160 or the like. Using transfer functions or similar techniques, the relative phases of all channels may be compared with one another at all selected frequencies. By keeping the one or more reference probes fixed at their unchanging locations while moving or monitoring the other tri-axial vibration sensors, it can be shown that the entire machine can be mapped with regard to amplitude and relative phase. This can be shown to be true even when there are more measurement points than channels of data collection. With this information, an operating deflection shape can be created that can show dynamic movements of the machine in 3 D, which can provide an invaluable diagnostic tool. In example embodiments, the one or more reference probes can

provide relative phase, rather than absolute phase. It will be appreciated in light of the disclosure that relative phase may not be as valuable absolute phase for some purposes, but the relative phase the information can still be shown to be very useful.

[0319] In example embodiments, the sampling rates used during the vibration survey can be digitally synchronized to predetermined operational frequencies that can relate to pertinent parameters of the machine such as rotating or oscillating speed. Doing this, permits extracting even more information using synchronized averaging techniques. It will be appreciated in light of the disclosure that this can be done without the use of a key phasor or a reference pulse from a rotating shaft, which is usually not available for route collected data. As such, non-synchronous signals can be removed from a complex signal without the need to deploy synchronous averaging using the key phasor. This can be shown to be very powerful when analyzing a particular pinion in a gearbox or generally applied to any component within a complicated mechanical mechanism. In many instances, the key phasor or the reference pulse is rarely available with route collected data, but the techniques disclosed herein can overcome this absence. In example embodiments, there can be multiple shafts running at different speeds within the machine being analyzed. In certain instances, there can be a single-axis reference probe for each shaft. In other instances, it is possible to relate the phase of one shaft to another shaft using only one single axis reference probe on one shaft at its unchanging location. In example embodiments, variable speed equipment can be more readily analyzed with relatively longer duration of data relative to single speed equipment. The vibration survey can be conducted at several machine speeds within the same contiguous set of vibration data using the same techniques disclosed herein. These techniques can also permit the study of the change of the relationship between vibration and the change of the rate of speed that was not available before.

[0320] In example embodiments, there are numerous analytical techniques that can emerge because raw waveform data can be captured in a gap-free digital format as disclosed herein. The gap-free digital format can facilitate many paths to analyze the waveform data in many ways after the fact to identify specific problems. The vibration data collected in accordance with the techniques disclosed herein can provide the analysis of transient, semi-periodic and very low frequency phenomena. The waveform data acquired in accordance with the present disclosure can contain relatively longer streams of raw gap-free waveform data that can be conveniently played back as needed, and on which many and varied sophisticated analytical techniques can be performed. A large number of such techniques can provide for various forms of filtering to extract low amplitude modulations from transient impact data that can be included in the relatively longer stream of raw gap-free waveform data. It will be appreciated in light of the disclosure that in past data collection practices, these types of phenomena were typically lost by the averaging process of the spectral processing algorithms because the goal of the previous data acquisition module was purely periodic signals; or these phenomena were lost to file size reduction methodologies due to the fact that much of the content from an original raw signal was typically discarded knowing it would not be used.

[0321] In example embodiments, there is a method of monitoring vibration of a machine having at least one shaft supported by a set of bearings. The method includes monitoring a first data channel assigned to a single-axis sensor at an unchanging location associated with the machine. The method also includes monitoring a second, third, and fourth data channel assigned to a three-axis sensor. The method further includes recording gap-free digital waveform data simultaneously from all of the data channels while the machine is in operation; and determining a change in relative phase based on the digital waveform data. The method also includes the tri-axial sensor being located at a plurality of positions associated with the machine while obtaining the digital waveform. In example embodiments, the second, third, and fourth channels are assigned together to a sequence of tri-axial sensors each located at different positions associated with the machine. In example embodiments, the data is received from all of the sensors on all of their channels simultaneously.

[0322] The method also includes determining an operating deflection shape based on the change in relative phase information and the waveform data. In example embodiments, the unchanging location of the reference sensor is a position associated with a shaft of the machine. In example embodiments, the tri-axial sensors in the sequence of the tri-axial sensors are each located at different positions and are each associated with different bearings in the machine. In example embodiments, the unchanging location is a position associated with a shaft of the machine and, wherein, the tri-axial sensors in the sequence of the tri-axial sensors are each located at different positions and are each associated with different bearings that support the shaft in the machine. The various embodiments include methods of sequentially monitoring vibration or similar process parameters and signals of a rotating or oscillating machine or analogous process machinery from a number of channels simultaneously, which can be known as an ensemble. In various examples, the ensemble can include one to eight channels. In further examples, an ensemble can represent a logical measurement grouping on the equipment being monitored whether those measurement locations are temporary for measurement, supplied by the original equipment manufacturer, retrofit at a later date, or one or more combinations thereof.

[0323] In one example, an ensemble can monitor bearing vibration in a single direction. In a further example, an ensemble can monitor three different directions (e.g., orthogonal directions) using a tri-axial sensor. In yet further examples, an ensemble can monitor four or more channels where the first channel can monitor a single axis vibration sensor, and the second, the third, and the fourth channels can monitor each of the three directions of the tri-axial sensor. In other examples, the ensemble can be fixed to a group of adjacent bearings on the same piece of equipment or an associated shaft. The various embodiments provide methods that include strategies for collecting waveform data from various ensembles deployed in vibration studies or the like in a relatively more efficient manner. The methods also include simultaneously monitoring of a reference channel assigned to an unchanging reference location associated with the ensemble monitoring the machine. The cooperation with the reference channel can be shown to support a more complete correlation of the collected waveforms from the ensembles. The reference sensor on the reference channel can be a single axis vibration sensor, or a phase reference sensor that can be triggered by a reference location on a rotating shaft or the like. As disclosed herein, the methods can further

include recording gap-free digital waveform data simultaneously from all of the channels of each ensemble at a relatively high rate of sampling so as to include all frequencies deemed necessary for the proper analysis of the machinery being monitored while it is in operation. The data from the ensembles can be streamed gap-free to a storage medium for subsequent processing that can be connected to a cloud network facility, a local data link, BLUETOOTH™ connectivity, cellular data connectivity, or the like.

[0324] In example embodiments, the methods disclosed herein include strategies for collecting data from the various ensembles including digital signal processing techniques that can be subsequently applied to data from the ensembles to emphasize or better isolate specific frequencies or waveform phenomena. This can be in contrast with current methods that collect multiple sets of data at different sampling rates, or with different hardware filtering configurations including integration, that provide relatively less post-processing flexibility because of the commitment to these same (known as a priori hardware configurations). These same hardware configurations can also be shown to increase time of the vibration survey due to the latency delays associated with configuring the hardware for each independent test. In example embodiments, the methods for collecting data from various ensembles include data marker technology that can be used for classifying sections of streamed data as homogenous and belonging to a specific ensemble. In one example, a classification can be defined as operating speed. In doing so, a multitude of ensembles can be created from what conventional systems would collect as only one. The many embodiments include post-processing analytic techniques for comparing the relative phases of all the frequencies of interest not only between each channel of the collected ensemble but also between all of the channels of all of the ensembles being monitored, when applicable.

[0325] The present disclosure can include markers that can be applied to a time mark or a sample length within the raw waveform data. The markers generally fall into two categories: preset or dynamic. The preset markers can correlate to preset or existing operating conditions (e.g., load, head pressure, air flow cubic feet per minute, ambient temperature, RPMs, and the like.). These preset markers can be fed into the data acquisition system directly. In certain instances, the preset markers can be collected on data channels in parallel with the waveform data (e.g., waveforms for vibration, current, voltage, etc.). Alternatively, the values for the preset markers can be entered manually.

[0326] For dynamic markers such as trending data, it can be important to compare similar data like comparing vibration amplitudes and patterns with a repeatable set of operating parameters. One example of the present disclosure includes one of the parallel channel inputs being a key phasor trigger pulse from an operating shaft that can provide RPM information at the instantaneous time of collection. In this example of dynamic markers, sections of collected waveform data can be marked with appropriate speeds or speed ranges.

[0327] The present disclosure can also include dynamic markers that can correlate to data that can be derived from post processing and analytics performed on the sample waveform. In further embodiments, the dynamic markers can also correlate to post-collection derived parameters including RPMs, as well as other operationally derived metrics such as alarm conditions like a

maximum RPM. In certain examples, many modern pieces of equipment that are candidates for a vibration survey with the portable data collection systems described herein do not include tachometer information. This can be true because it is not always practical or cost-justifiable to add a tachometer even though the measurement of RPM can be of primary importance for the vibration survey and analysis. It will be appreciated that for fixed speed machinery obtaining an accurate RPM measurement can be less important especially when the approximate speed of the machine can be ascertained before-hand; however, variable-speed drives are becoming increasingly prevalent. It will also be appreciated in light of the disclosure that various signal processing techniques can permit the derivation of RPM from the raw data without the need for a dedicated tachometer signal.

[0328] In many embodiments, the RPM information can be used to mark segments of the raw waveform data over its collection history. Further embodiments include techniques for collecting instrument data following a prescribed route of a vibration study. The dynamic markers can enable analysis and trending software to utilize multiple segments of the collection interval indicated by the markers (e.g., two minutes) as multiple historical collection ensembles, rather than just one as done in previous systems where route collection systems would historically store data for only one RPM setting. This could, in turn, be extended to any other operational parameter such as load setting, ambient temperature, and the like, as previously described. The dynamic markers, however, that can be placed in a type of index file pointing to the raw data stream can classify portions of the stream in homogenous entities that can be more readily compared to previously collected portions of the raw data stream.

[0329] The many embodiments include the hybrid relational metadata-binary storage approach that can use the best of pre-existing technologies for both relational and raw data streams. In example embodiments, the hybrid relational metadata - binary storage approach can marry them together with a variety of marker linkages. The marker linkages can permit rapid searches through the relational metadata and can allow for more efficient analyses of the raw data using conventional SQL techniques with pre-existing technology. This can be shown to permit utilization of many of the capabilities, linkages, compatibilities, and extensions that conventional database technologies do not provide.

[0330] The marker linkages can also permit rapid and efficient storage of the raw data using conventional binary storage and data compression techniques. This can be shown to permit utilization of many of the capabilities, linkages, compatibilities, and extensions that conventional raw data technologies provide such as TMDS (National Instruments), UFF (Universal File Format such as UFF58), and the like. The marker linkages can further permit using the marker technology links where a vastly richer set of data from the ensembles can be amassed in the same collection time as more conventional systems. The richer set of data from the ensembles can store data snapshots associated with predetermined collection criterion and the proposed system can derive multiple snapshots from the collected data streams utilizing the marker technology. In doing so, it can be shown that a relatively richer analysis of the collected data can be achieved. One such benefit can include more trending points of vibration at a specific frequency or order of running

speed versus RPM, load, operating temperature, flow rates, and the like, which can be collected for a similar time relative to what is spent collecting data with a conventional system.

[0331] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from machines, elements of the machines and the environment of the machines including heavy duty machines deployed at a local job site or at distributed job sites under common control. The heavy-duty machines may include earthmoving equipment, heavy duty on-road industrial vehicles, heavy duty off-road industrial vehicles, industrial machines deployed in various settings such as turbines, turbomachinery, generators, pumps, pulley systems, manifold and valve systems, and the like. In example embodiments, heavy industrial machinery may also include earth-moving equipment, earth-compacting equipment, hauling equipment, hoisting equipment, conveying equipment, aggregate production equipment, equipment used in concrete construction, and piledriving equipment. In examples, earth moving equipment may include excavators, backhoes, loaders, bulldozers, skid steer loaders, trenchers, motor graders, motor scrapers, crawler loaders, and wheeled loading shovels. In examples, construction vehicles may include dumpers, tankers, tippers, and trailers. In examples, material handling equipment may include cranes, conveyors, forklift, and hoists. In examples, construction equipment may include tunnel and handling equipment, road rollers, concrete mixers, hot mix plants, road making machines (compactors), stone crashers, pavers, slurry seal machines, spraying and plastering machines, and heavy-duty pumps. Further examples of heavy industrial equipment may include different systems such as implement traction, structure, power train, control, and information. Heavy industrial equipment may include many different powertrains and combinations thereof to provide power for locomotion and to also provide power to accessories and onboard functionality. In each of these examples, the platform 100 may deploy the local data collection system 102 into the environment 104 in which these machines, motors, pumps, and the like, operate and directly connected integrated into each of the machines, motors, pumps, and the like.

[0332] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from machines in operation and machines in being constructed such as turbine and generator sets like Siemens™ SGT6-5000F™ gas turbine, an SST-900™ steam turbine, an SGen6-1000A™ generator, and an SGen6-100A™ generator, and the like. In example embodiments, the local data collection system 102 may be deployed to monitor steam turbines as they rotate in the currents caused by hot water vapor that may be directed through the turbine but otherwise generated from a different source such as from gas-fired burners, nuclear cores, molten salt loops and the like. In these systems, the local data collection system 102 may monitor the turbines and the water or other fluids in a closed loop cycle in which water condenses and is then heated until it evaporates again. The local data collection system 102 may monitor the steam turbines separately from the fuel source deployed to heat the water to steam. In examples, working temperatures of steam turbines may be between 500 and 650 °C. In many embodiments, an array of steam turbines may be arranged and configured for high, medium, and low pressure, so they may optimally convert the respective steam pressure into rotational movement.

[0333] The local data collection system 102 may also be deployed in a gas turbines arrangement and therefore not only monitor the turbine in operation but also monitor the hot combustion gases feed into the turbine that may be in excess of 1,500 °C. Because these gases are much hotter than those in steam turbines, the blades may be cooled with air that may flow out of small openings to create a protective film or boundary layer between the exhaust gases and the blades. This temperature profile may be monitored by the local data collection system 102. Gas turbine engines, unlike typical steam turbines, include a compressor, a combustion chamber, and a turbine all of which are journaled for rotation with a rotating shaft. The construction and operation of each of these components may be monitored by the local data collection system 102.

[0334] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from water turbines serving as rotary engines that may harvest energy from moving water and are used for electric power generation. The type of water turbine or hydro-power selected for a project may be based on the height of standing water, often referred to as head, and the flow (or volume of water) at the site. In this example, a generator may be placed at the top of a shaft that connects to the water turbine. As the turbine catches the naturally moving water in its blade and rotates, the turbine sends rotational power to the generator to generate electrical energy. In doing so, the platform 100 may monitor signals from the generators, the turbines, the local water system, flow controls such as dam windows and sluices. Moreover, the platform 100 may monitor local conditions on the electric grid including load, predicted demand, frequency response, and the like, and include such information in the monitoring and control deployed by platform 100 in these hydroelectric settings.

[0335] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from energy production environments, such as thermal, nuclear, geothermal, chemical, biomass, carbon-based fuels, hybrid-renewable energy plants, and the like. Many of these plants may use multiple forms of energy harvesting equipment like wind turbines, hydro turbines, and steam turbines powered by heat from nuclear, gas-fired, solar, and molten salt heat sources. In example embodiments, elements in such systems may include transmission lines, heat exchangers, desulphurization scrubbers, pumps, coolers, recuperators, chillers, and the like. In example embodiments, certain implementations of turbomachinery, turbines, scroll compressors, and the like may be configured in arrayed control so as to monitor large facilities creating electricity for consumption, providing refrigeration, creating steam for local manufacture and heating, and the like, and that arrayed control platforms may be provided by the provider of the industrial equipment such as Honeywell and their Experion™ PKS platform. In example embodiments, the platform 100 may specifically communicate with and integrate the local manufacturer-specific controls and may allow equipment from one manufacturer to communicate with other equipment. Moreover, the platform 100 provides allows for the local data collection system 102 to collect information across systems from many different manufacturers. In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from marine industrial equipment,

marine diesel engines, shipbuilding, oil and gas plants, refineries, petrochemical plant, ballast water treatment solutions, marine pumps and turbines, and the like.

[0336] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from heavy industrial equipment and processes including monitoring one or more sensors. By way of this example, sensors may be devices that may be used to detect or respond to some type of input from a physical environment, such as an electrical, heat, or optical signal. In example embodiments, the local data collection system 102 may include multiple sensors such as, without limitation, a temperature sensor, a pressure sensor, a torque sensor, a flow sensor, a heat sensor, a smoke sensor, an arc sensor, a radiation sensor, a position sensor, an acceleration sensor, a strain sensor, a pressure cycle sensor, a pressure sensor, an air temperature sensor, and the like. The torque sensor may encompass a magnetic twist angle sensor. In one example, the torque and speed sensors in the local data collection system 102 may be similar to those discussed in U.S. Patent Number 8,352,149 to Meachem, issued 8 January 2013 and hereby incorporated by reference as if fully set forth herein. In example embodiments, one or more sensors may be provided such as a tactile sensor, a biosensor, a chemical sensor, an image sensor, a humidity sensor, an inertial sensor, and the like.

[0337] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from sensors that may provide signals for fault detection including excessive vibration, incorrect material, incorrect material properties, trueness to the proper size, trueness to the proper shape, proper weight, trueness to balance. Additional fault sensors include those for inventory control and for inspections such as to confirm that parts are packaged to plan, parts are to tolerance in a plan, occurrence of packaging damage or stress, and sensors that may indicate the occurrence of shock or damage in transit. Additional fault sensors may include detection of the lack of lubrication, over lubrication, the need for cleaning of the sensor detection window, the need for maintenance due to low lubrication, the need for maintenance due to blocking or reduced flow in a lubrication region, and the like.

[0338] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 that includes aircraft operations and manufacture including monitoring signals from sensors for specialized applications such as sensors used in an aircraft's Attitude and Heading Reference System (AHRS), such as gyroscopes, accelerometers, and magnetometers. In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from image sensors such as semiconductor charge coupled devices (CCDs), active pixel sensors, in complementary metal–oxide–semiconductor (CMOS) or N-type metal-oxide-semiconductor (NMOS, Live MOS) technologies. In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from sensors such as an infra-red (IR) sensor, an ultraviolet (UV) sensor, a touch sensor, a proximity sensor, and the like. In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from sensors configured for optical character recognition (OCR), reading barcodes, detecting surface acoustic waves, detecting transponders,

communicating with home automation systems, medical diagnostics, health monitoring, and the like.

[0339] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from sensors such as a Micro-Electro-Mechanical Systems (MEMS) sensor, such as ST Microelectronic's™ LSM303AH smart MEMS sensor, which may include an ultra-low-power high-performance system-in-package featuring a 3D digital linear acceleration sensor and a 3D digital magnetic sensor.

[0340] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from additional large machines such as turbines, windmills, industrial vehicles, robots, and the like. These large mechanical machines include multiple components and elements providing multiple subsystems on each machine. To that end, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from individual elements such as axles, bearings, belts, buckets, gears, shafts, gear boxes, cams, carriages, camshafts, clutches, brakes, drums, dynamos, feeds, flywheels, gaskets, pumps, jaws, robotic arms, seals, sockets, sleeves, valves, wheels, actuators, motors, servomotor, and the like. Many of the machines and their elements may include servomotors. The local data collection system 102 may monitor the motor, the rotary encoder, and the potentiometer of the servomechanism to provide three-dimensional detail of position, placement, and progress of industrial processes.

[0341] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from gear drives, powertrains, transfer cases, multispeed axles, transmissions, direct drives, chain drives, belt-drives, shaft-drives, magnetic drives, and similar meshing mechanical drives. In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from fault conditions of industrial machines that may include overheating, noise, grinding gears, locked gears, excessive vibration, wobbling, under-inflation, over-inflation, and the like. Operation faults, maintenance indicators, and interactions from other machines may cause maintenance or operational issues may occur during operation, during installation, and during maintenance. The faults may occur in the mechanisms of the industrial machines but may also occur in infrastructure that supports the machine such as its wiring and local installation platforms. In example embodiments, the large industrial machines may face different types of fault conditions such as overheating, noise, grinding gears, excessive vibration of machine parts, fan vibration problems, problems with large industrial machines rotating parts.

[0342] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor signals from industrial machinery including failures that may be caused by premature bearing failure that may occur due to contamination or loss of bearing lubricant. In another example, a mechanical defect such as misalignment of bearings may occur. Many factors may contribute to the failure such as metal fatigue, therefore, the local data collection system 102 may monitor cycles and local stresses. By way of this example, the platform 100 may monitor the incorrect operation of machine parts, lack of maintenance and

servicing of parts, corrosion of vital machine parts, such as couplings or gearboxes, misalignment of machine parts, and the like. Though the fault occurrences cannot be completely stopped, many industrial breakdowns may be mitigated to reduce operational and financial losses. The platform 100 provides real-time monitoring and predictive maintenance in many industrial environments wherein it has been shown to present a cost-savings over regularly-scheduled maintenance processes that replace parts according to a rigid expiration of time and not actual load and wear and tear on the element or machine. To that end, the platform 10 may provide reminders of, or perform some, preventive measures such as adhering to operating manual and mode instructions for machines, proper lubrication, and maintenance of machine parts, minimizing or eliminating overrun of machines beyond their defined capacities, replacement of worn but still functional parts as needed, properly training the personnel for machine use, and the like.

[0343] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 to monitor multiple signals that may be carried by a plurality of physical, electronic, and symbolic formats or signals. The platform 100 may employ signal processing including a plurality of mathematical, statistical, computational, heuristic, and linguistic representations and processing of signals and a plurality of operations needed for extraction of useful information from signal processing operations such as techniques for representation, modeling, analysis, synthesis, sensing, acquisition, and extraction of information from signals. In examples, signal processing may be performed using a plurality of techniques, including but not limited to transformations, spectral estimations, statistical operations, probabilistic and stochastic operations, numerical theory analysis, data mining, and the like. The processing of various types of signals forms the basis of many electrical or computational processes. As a result, signal processing applies to almost all disciplines and applications in the industrial environment such as audio and video processing, image processing, wireless communications, process control, industrial automation, financial systems, feature extraction, quality improvements such as noise reduction, image enhancement, and the like. Signal processing for images may include pattern recognition for manufacturing inspections, quality inspection, and automated operational inspection and maintenance. The platform 100 may employ many pattern recognition techniques including those that may classify input data into classes based on key features with the objective of recognizing patterns or regularities in data. The platform 100 may also implement pattern recognition processes with machine learning operations and may be used in applications such as computer vision, speech and text processing, radar processing, handwriting recognition, CAD systems, and the like. The platform 100 may employ supervised classification and unsupervised classification. The supervised learning classification algorithms may be based to create classifiers for image or pattern recognition, based on training data obtained from different object classes. The unsupervised learning classification algorithms may operate by finding hidden structures in unlabeled data using advanced analysis techniques such as segmentation and clustering. For example, some of the analysis techniques used in unsupervised learning may include K-means clustering, Gaussian mixture models, Hidden Markov models, and the like. The algorithms used in supervised and unsupervised learning methods of pattern recognition enable the use of pattern

recognition in various high precision applications. The platform 100 may use pattern recognition in face detection related applications such as security systems, tracking, sports related applications, fingerprint analysis, medical and forensic applications, navigation and guidance systems, vehicle tracking, public infrastructure systems such as transport systems, license plate monitoring, and the like.

[0344] In example embodiments, the platform 100 may include the local data collection system 102 deployed in the environment 104 using machine learning to enable derivation-based learning outcomes from computers without the need to program them. The platform 100 may, therefore, learn from and make decisions on a set of data, by making data-driven predictions and adapting according to the set of data. In example embodiments, machine learning may involve performing a plurality of machine learning tasks by machine learning systems, such as supervised learning, unsupervised learning, and reinforcement learning. Supervised learning may include presenting a set of example inputs and desired outputs to the machine learning systems. Unsupervised learning may include the learning algorithm itself structuring its input by methods such as pattern detection and/or feature learning. Reinforcement learning may include the machine learning systems performing in a dynamic environment and then providing feedback about correct and incorrect decisions. In examples, machine learning may include a plurality of other tasks based on an output of the machine learning system. In examples, the tasks may also be classified as machine learning problems such as classification, regression, clustering, density estimation, dimensionality reduction, anomaly detection, and the like. In examples, machine learning may include a plurality of mathematical and statistical techniques. In examples, the many types of machine learning algorithms may include decision tree based learning, association rule learning, deep learning, artificial neural networks, genetic learning algorithms, inductive logic programming, support vector machines (SVMs), Bayesian network, reinforcement learning, representation learning, rule-based machine learning, sparse dictionary learning, similarity and metric learning, learning classifier systems (LCS), logistic regression, random forest, K-Means, gradient boost and adaboost, K-nearest neighbors (KNN), a priori algorithms, and the like. In example embodiments, certain machine learning algorithms may be used (such as genetic algorithms defined for solving both constrained and unconstrained optimization problems that may be based on natural selection, the process that drives biological evolution). By way of this example, genetic algorithms may be deployed to solve a variety of optimization problems that are not well suited for standard optimization algorithms, including problems in which the objective functions are discontinuous, not differentiable, stochastic, or highly nonlinear. In an example, the genetic algorithm may be used to address problems of mixed integer programming, where some components restricted to being integer-valued. Genetic algorithms and machine learning techniques and systems may be used in computational intelligence systems, computer vision, Natural Language Processing (NLP), recommender systems, reinforcement learning, building graphical models, and the like. By way of this example, the machine learning systems may be used to perform intelligent computing based control and be responsive to tasks in a wide variety of systems (such as interactive websites and portals, brain-machine interfaces, online security and fraud detection systems, medical applications

such as diagnosis and therapy assistance systems, classification of DNA sequences, and the like). In examples, machine learning systems may be used in advanced computing applications (such as online advertising, natural language processing, robotics, search engines, software engineering, speech and handwriting recognition, pattern matching, game playing, computational anatomy, bioinformatics systems and the like). In an example, machine learning may also be used in financial and marketing systems (such as for user behavior analytics, online advertising, economic estimations, financial market analysis, and the like).

[0345] Additional details are provided below in connection with the methods, systems, devices, and components depicted in connection with Figures 1 through 6. In example embodiments, methods and systems are disclosed herein for cloud-based, machine pattern recognition based on fusion of remote, analog industrial sensors. For example, data streams from vibration, pressure, temperature, accelerometer, magnetic, electrical field, and other analog sensors may be multiplexed or otherwise fused, relayed over a network, and fed into a cloud-based machine learning facility, which may employ one or more models relating to an operating characteristic of an industrial machine, an industrial process, or a component or element thereof. A model may be created by a human who has experience with the industrial environment and may be associated with a training data set (such as models created by human analysis or machine analysis of data that is collected by the sensors in the environment, or sensors in other similar environments. The learning machine may then operate on other data, initially using a set of rules or elements of a model, such as to provide a variety of outputs, such as classification of data into types, recognition of certain patterns (such as those indicating the presence of faults, orthoses indicating operating conditions, such as fuel efficiency, energy production, or the like). The machine learning facility may take feedback, such as one or more inputs or measures of success, such that it may train, or improve, its initial model (such as improvements by adjusting weights, rules, parameters, or the like, based on the feedback). For example, a model of fuel consumption by an industrial machine may include physical model parameters that characterize weights, motion, resistance, momentum, inertia, acceleration, and other factors that indicate consumption, and chemical model parameters (such as those that predict energy produced and/or consumed e.g., such as through combustion, through chemical reactions in battery charging and discharging, and the like). The model may be refined by feeding in data from sensors disposed in the environment of a machine, in the machine, and the like, as well as data indicating actual fuel consumption, so that the machine can provide increasingly accurate, sensor-based, estimates of fuel consumption and can also provide output that indicate what changes can be made to increase fuel consumption (such as changing operation parameters of the machine or changing other elements of the environment, such as the ambient temperature, the operation of a nearby machine, or the like). For example, if a resonance effect between two machines is adversely affecting one of them, the model may account for this and automatically provide an output that results in changing the operation of one of the machines (such as to reduce the resonance, to increase fuel efficiency of one or both machines). By continuously adjusting parameters to cause outputs to match actual conditions, the machine learning facility may self-organize to provide a highly accurate model of the conditions of an environment (such as for

predicting faults, optimizing operational parameters, and the like). This may be used to increase fuel efficiency, to reduce wear, to increase output, to increase operating life, to avoid fault conditions, and for many other purposes.

[0346] Figure 10 illustrates components and interactions of a data collection architecture involving the application of cognitive and machine learning systems to data collection and processing. Referring to Figure 10, a data collection system 102 may be disposed in an environment (such as an industrial environment where one or more complex systems, such as electro-mechanical systems and machines are manufactured, assembled, or operated). The data collection system 102 may include onboard sensors and may take input, such as through one or more input interfaces or ports 4008, from one or more sensors (such as analog or digital sensors of any type disclosed herein) and from one or more input sources 116 (such as sources that may be available through Wi-Fi, BLUETOOTH™, NFC, or other local network connections or over the Internet). Sensors may be combined and multiplexed (such as with one or more multiplexers 4002). Data may be cached or buffered in a cache/buffer 4022 and made available to external systems, such as a remote host processing system 112 as described elsewhere in this disclosure (which may include an extensive processing architecture 4024, including any of the elements described in connection with other embodiments described throughout this disclosure and in the Figure), though one or more output interfaces and ports 4010 (which may in embodiments be separate from or the same as the input interfaces and ports 4008). The data collection system 102 may be configured to take input from a host processing system 112, such as input from an analytic system 4018, which may operate on data from the data collection system 102 and data from other input sources 116 to provide analytic results, which in turn may be provided as a learning feedback input 4012 to the data collection system, such as to assist in configuration and operation of the data collection system 102.

[0347] Combination of inputs (including selection of what sensors or input sources to turn "on" or "off") may be performed under the control of machine-based intelligence, such as using a local cognitive input selection system 4004, an optionally remote cognitive input selection system 4114, or a combination of the two. The cognitive input selection systems 4004, 4014 may use intelligence and machine learning capabilities described elsewhere in this disclosure, such as using detected conditions (such as conditions informed by the input sources 116 or sensors), state information (including state information determined by a machine state recognition system 4020 that may determine a state), such as relating to an operational state, an environmental state, a state within a known process or workflow, a state involving a fault or diagnostic condition, or many others. This may include optimization of input selection and configuration based on learning feedback from the learning feedback system 4012, which may include providing training data (such as from the host processing system 112 or from other data collection systems 102 either directly or from the host 112) and may include providing feedback metrics, such as success metrics calculated within the analytic system 4018 of the host processing system 112. For example, if a data stream consisting of a particular combination of sensors and inputs yields positive results in a given set of conditions (such as providing improved pattern recognition, improved prediction, improved diagnosis, improved yield, improved return on investment, improved efficiency, or the like), then metrics

relating to such results from the analytic system 4018 can be provided via the learning feedback system 4012 to the cognitive input selection systems 4004, 4014 to help configure future data collection to select that combination in those conditions (allowing other input sources to be de-selected, such as by powering down the other sensors). In example embodiments, selection and de-selection of sensor combinations, under control of one or more of the cognitive input selection systems 4004, may occur with automated variation, such as using genetic programming techniques, based on learning feedback 4012, such as from the analytic system 4018, effective combinations for a given state or set of conditions are promoted, and less effective combinations are demoted, resulting in progressive optimization and adaptation of the local data collection system to each unique environment. Thus, an automatically adapting, multi-sensor data collection system is provided, where cognitive input selection is used (with feedback) to improve the effectiveness, efficiency, or other performance parameters of the data collection system within its particular environment. Performance parameters may relate to overall system metrics (such as financial yields, process optimization results, energy production or usage, and the like), analytic metrics (such as success in recognizing patterns, making predictions, classifying data, or the like), and local system metrics (such as bandwidth utilization, storage utilization, power consumption, and the like). In example embodiments, the analytic system 4018, the state system 4020 and the cognitive input selection system 4114 of a host may take data from multiple data collection systems 102, such that optimization (including of input selection) may be undertaken through coordinated operation of multiple systems 102. For example, the cognitive input selection system 4114 may understand that if one data collection system 102 is already collecting vibration data for an X-axis, the X-axis vibration sensor for the other data collection system might be turned off, in favor of getting Y-axis data from the other data collector 102. Thus, through coordinated collection by the host cognitive input selection system 4114, the activity of multiple collectors 102, across a host of different sensors, can provide for a rich data set for the host processing system 112, without wasting energy, bandwidth, storage space, or the like. As noted above, optimization may be based on overall system success metrics, analytic success metrics, and local system metrics, or a combination of the above.

[0348] Methods and systems are disclosed herein for cloud-based, machine pattern analysis of state information from multiple industrial sensors to provide anticipated state information for an industrial system. In example embodiments, machine learning may take advantage of a state machine, such as tracking states of multiple analog and/or digital sensors, feeding the states into a pattern analysis facility, and determining anticipated states of the industrial system based on historical data about sequences of state information. For example, where a temperature state of an industrial machine exceeds a certain threshold and is followed by a fault condition, such as breaking down of a set of bearings, that temperature state may be tracked by a pattern recognizer, which may produce an output data structure indicating an anticipated bearing fault state (whenever an input state of a high temperature is recognized). A wide range of measurement values and anticipated states may be managed by a state machine, relating to temperature, pressure, vibration, acceleration, momentum, inertia, friction, heat, heat flux, galvanic states, magnetic field states,

electrical field states, capacitance states, charge and discharge states, motion, position, and many others. States may comprise combined states, where a data structure includes a series of states, each of which is represented by a place in a byte-like data structure. For example, an industrial machine may be characterized by a genetic structure, such as one that provides pressure, temperature, vibration, and acoustic data, the measurement of which takes one place in the data structure, so that the combined state can be operated on as a byte-like structure, such as a structure for compactly characterizing the current combined state of the machine or environment, or compactly characterizing the anticipated state. This byte-like structure can be used by a state machine for machine learning, such as pattern recognition that operates on the structure to determine patterns that reflect combined effects of multiple conditions. A wide variety of such structures can be tracked and used, such as in machine learning, representing various combinations, of various length, of the different elements that can be sensed in an industrial environment. In example embodiments, byte-like structures can be used in a genetic programming technique, such as by substituting different types of data, or data from varying sources, and tracking outcomes over time, so that one or more favorable structures emerges based on the success of those structures when used in real world situations, such as indicating successful predictions of anticipated states, or achievement of success operational outcomes, such as increased efficiency, successful routing of information, achieving increased profits, or the like. That is, by varying what data types and sources are used in byte-like structures that are used for machine optimization over time, a genetic programming-based machine learning facility can "evolve" a set of data structures, consisting of a favorable mix of data types (e.g., pressure, temperature, and vibration), from a favorable mix of data sources (e.g., temperature is derived from sensor X, while vibration comes from sensor Y), for a given purpose. Different desired outcomes may result in different data structures that are best adapted to support effective achievement of those outcomes over time with application of machine learning and promotion of structures with favorable results for the desired outcome in question by genetic programming. The promoted data structures may provide compact, efficient data for various activities as described throughout this disclosure, including being stored in data pools (which may be optimized by storing favorable data structures that provide the best operational results for a given environment), being presented in data marketplaces (such as being presented as the most effective structures for a given purpose), and the like.

[0349] In example embodiments, a platform is provided having cloud-based, machine pattern analysis of state information from multiple analog industrial sensors to provide anticipated state information for an industrial system. In example embodiments, the host processing system 112, such as disposed in the cloud, may include the state system 4020, which may be used to infer or calculate a current state or to determine an anticipated future state relating to the data collection system 102 or some aspect of the environment in which the data collection system 102 is disposed, such as the state of a machine, a component, a workflow, a process, an event (e.g., whether the event has occurred), an object, a person, a condition, a function, or the like. Maintaining state information allows the host processing system 112 to undertake analysis, such as in one or more analytic systems 4018, to determine contextual information, to apply semantic and conditional

logic, and perform many other functions as enabled by the processing architecture 4024 described throughout this disclosure.

[0350] In example embodiments, a platform is provided having cloud-based policy automation engine for IoT, with creation, deployment, and management of IoT devices. In example embodiments, the platform 100 includes (or is integrated with, or included in) the host processing system 112, such as on a cloud platform, a policy automation engine 4032 for automating creation, deployment, and management of policies to IoT devices. Polices, which may include access policies, network usage policies, storage usage policies, bandwidth usage policies, device connection policies, security policies, rule-based policies, role-based polices, and others, may be required to govern the use of IoT devices. For example, as IoT devices may have many different network and data communications to other devices, policies may be needed to indicate to what devices a given device can connect, what data can be passed on, and what data can be received. As billions of devices with countless potential connections are expected to be deployed in the near future, it becomes impossible for humans to configure policies for IoT devices on a connection-by-connection basis. Accordingly, an intelligent policy automation engine 4032 may include cognitive features for creating, configuring, and managing policies. The policy automation engine 4032 may consume information about possible policies, such as from a policy database or library, which may include one or more public sources of available policies. These may be written in one or more conventional policy languages or scripts. The policy automation engine 4032 may apply the policies according to one or more models, such as based on the characteristics of a given device, machine, or environment. For example, a large machine, such as a machine for power generation, may include a policy that only a verifiably local controller can change certain parameters of the power generation, thereby avoiding a remote "takeover" by a hacker. This may be accomplished in turn by automatically finding and applying security policies that bar connection of the control infrastructure of the machine to the Internet, by requiring access authentication, or the like. The policy automation engine 4032 may include cognitive features, such as varying the application of policies, the configuration of policies, and the like (such as features based on state information from the state system 4020). The policy automation engine 4032 may take feedback, as from the learning feedback system 4012, such as based on one or more analytic results from the analytic system 4018, such as based on overall system results (such as the extent of security breaches, policy violations, and the like), local results, and analytic results. By variation and selection based on such feedback, the policy automation engine 4032 can, over time, learn to automatically create, deploy, configure, and manage policies across very large numbers of devices, such as managing policies for configuration of connections among IoT devices.

[0351] Methods and systems are disclosed herein for on-device sensor fusion and data storage for industrial IoT devices, including on-device sensor fusion and data storage for an industrial IoT device, where data from multiple sensors is multiplexed at the device for storage of a fused data stream. For example, pressure and temperature data may be multiplexed into a data stream that combines pressure and temperature in a time series, such as in a byte-like structure (where time, pressure, and temperature are bytes in a data structure, so that pressure and temperature remain

linked in time, without requiring separate processing of the streams by outside systems), or by adding, dividing, multiplying, subtracting, or the like, such that the fused data can be stored on the device. Any of the sensor data types described throughout this disclosure can be fused in this manner and stored in a local data pool, in storage, or on an IoT device, such as a data collector, a component of a machine, or the like.

[0352] In example embodiments, a platform is provided having on-device sensor fusion and data storage for industrial IoT devices. In example embodiments, a cognitive system is used for a self-organizing storage system 4028 for the data collection system 102. Sensor data, and in particular analog sensor data, can consume large amounts of storage capacity, in particular where a data collector 102 has multiple sensor inputs onboard or from the local environment. Simply storing all the data indefinitely is not typically a favorable option, and even transmitting all of the data may strain bandwidth limitations, exceed bandwidth permissions (such as exceeding cellular data plan capacity), or the like. Accordingly, storage strategies are needed. These typically include capturing only portions of the data (such as snapshots), storing data for limited time periods, storing portions of the data (such as intermediate or abstracted forms), and the like. With many possible selections among these and other options, determining the correct storage strategy may be highly complex. In example embodiments, the self-organizing storage system 4028 may use a cognitive system, based on learning feedback 4012, and use various metrics from the analytic system 4018 or other system of the host cognitive input selection system 4114, such as overall system metrics, analytic metrics, and local performance indicators. The self-organizing storage system 4028 may automatically vary storage parameters, such as storage locations (including local storage on the data collection system 102, storage on nearby data collection systems 102 (such as using peer-to-peer organization) and remote storage, such as network-based storage), storage amounts, storage duration, type of data stored (including individual sensors or input sources 116, as well as various combined or multiplexed data, such as selected under the cognitive input selection systems 4004, 4014), storage type (such as using RAM, Flash, or other short-term memory versus available hard drive space), storage organization (such as in raw form, in hierarchies, and the like), and others. Variation of the parameters may be undertaken with feedback, so that over time the data collection system 102 adapts its storage of data to optimize itself to the conditions of its environment, such as a particular industrial environment, in a way that results in it storing the data that is needed in the right amounts and of the right type for availability to users.

[0353] In example embodiments, the local cognitive input selection system 4004 may organize fusion of data for various onboard sensors, external sensors (such as in the local environment) and other input sources 116 to the local collection system 102 into one or more fused data streams, such as using the multiplexer 4002 to create various signals that represent combinations, permutations, mixes, layers, abstractions, data-metadata combinations, and the like of the source analog and/or digital data that is handled by the data collection system 102. The selection of a particular fusion of sensors may be determined locally by the cognitive input selection system 4004, such as based on learning feedback from the learning feedback system 4012, such as various overall system, analytic system and local system results and metrics. In example embodiments, the

system may learn to fuse particular combinations and permutations of sensors, such as in order to best achieve correct anticipation of state, as indicated by feedback of the analytic system 4018 regarding its ability to predict future states, such as the various states handled by the state system 4020. For example, the input selection system 4004 may indicate selection of a sub-set of sensors among a larger set of available sensors, and the inputs from the selected sensors may be combined, such as by placing input from each of them into a byte of a defined, multi-bit data structure (such as a combination by taking a signal from each at a given sampling rate or time and placing the result into the byte structure, then collecting and processing the bytes over time), by multiplexing in the multiplexer 4002, such as a combination by additive mixing of continuous signals, and the like. Any of a wide range of signal processing and data processing techniques for combination and fusing may be used, including convolutional techniques, coercion techniques, transformation techniques, and the like. The particular fusion in question may be adapted to a given situation by cognitive learning, such as by having the cognitive input selection system 4004 learn, based on feedback 4012 from results (such as feedback conveyed by the analytic system 4018), such that the local data collection system 102 executes context-adaptive sensor fusion.

[0354] In example embodiments, the analytic system 4018 may apply to any of a wide range of analytic techniques, including statistical and econometric techniques (such as linear regression analysis, use similarity matrices, heat map based techniques, and the like), reasoning techniques (such as Bayesian reasoning, rule-based reasoning, inductive reasoning, and the like), iterative techniques (such as feedback, recursion, feed-forward and other techniques), signal processing techniques (such as Fourier and other transforms), pattern recognition techniques (such as Kalman and other filtering techniques), search techniques, probabilistic techniques (such as random walks, random forest algorithms, and the like), simulation techniques (such as random walks, random forest algorithms, linear optimization and the like), and others. This may include computation of various statistics or measures. In example embodiments, the analytic system 4018 may be disposed, at least in part, on a data collection system 102, such that a local analytic system can calculate one or more measures, such as measures relating to any of the items noted throughout this disclosure. For example, measures of efficiency, power utilization, storage utilization, redundancy, entropy, and other factors may be calculated onboard, so that the data collection 102 can enable various cognitive and learning functions noted throughout this disclosure without dependence on a remote (e.g., cloud-based) analytic system.

[0355] In example embodiments, the host processing system 112, a data collection system 102, or both, may include, connect to, or integrate with, a self-organizing networking system 4020, which may comprise a cognitive system for providing machine-based, intelligent or organization of network utilization for transport of data in a data collection system, such as for handling analog and other sensor data, or other source data, such as among one or more local data collection systems 102 and a host system 112. This may include organizing network utilization for source data delivered to data collection systems, for feedback data, such as analytic data provided to or via a learning feedback system 4012, data for supporting a marketplace (such as described in connection

with other embodiments), and output data provided via output interfaces and ports 4010 from one or more data collection systems 102.

[0356] Methods and systems are disclosed herein for a self-organizing data marketplace for industrial IoT data, including where available data elements are organized in the marketplace for consumption by consumers based on training a self-organizing facility with a training set and feedback from measures of marketplace success. A marketplace may be set up initially to make available data collected from one or more industrial environments, such as presenting data by type, by source, by environment, by machine, by one or more patterns, or the like (such as in a menu or hierarchy). The marketplace may vary the data collected, the organization of the data, the presentation of the data (including pushing the data to external sites, providing links, configuring APIs by which the data may be accessed, and the like), the pricing of the data, or the like, such as under machine learning, which may vary different parameters of any of the foregoing. The machine learning facility may manage all of these parameters by self-organization, such as by varying parameters over time (including by varying elements of the data types presented), the data sourced used to obtain each type of data, the data structures presented (such as byte-like structures, fused or multiplexed structures (such as representing multiple sensor types), and statistical structures (such as representing various mathematical products of sensor information), among others), the pricing for the data, where the data is presented, how the data is presented (such as by APIs, by links, by push messaging, and the like), how the data is stored, how the data is obtained, and the like. As parameters are varied, feedback may be obtained as to measures of success, such as number of views, yield (e.g., price paid) per access, total yield, per unit profit, aggregate profit, and many others, and the self-organizing machine learning facility may promote configurations that improve measures of success and demote configurations that do not, so that, over time, the marketplace is progressively configured to present favorable combinations of data types (e.g., those that provide robust prediction of anticipated states of particular industrial environments of a given type), from favorable sources (e.g., those that are reliable, accurate and low priced), with effective pricing (e.g., pricing that tends to provide high aggregate profit from the marketplace). The marketplace may include spiders, web crawlers, and the like to seek input data sources, such as finding data pools, connected IoT devices, and the like that publish potentially relevant data. These may be trained by human users and improved by machine learning in a manner similar to that described elsewhere in this disclosure.

[0357] In example embodiments, a platform is provided having a self-organizing data marketplace for industrial IoT data. Referring to Figure 11, in example embodiments, a platform is provided having a cognitive data marketplace 4102, referred to in some cases as a self-organizing data marketplace, for data collected by one or more data collection systems 102 or for data from other sensors or input sources 116 that are located in various data collection environments, such as industrial environments. In addition to data collection systems 102, this may include data collected, handled or exchanged by IoT devices, such as cameras, monitors, embedded sensors, mobile devices, diagnostic devices and systems, instrumentation systems, telematics systems, and the like, such as for monitoring various parameters and features of machines, devices, components, parts,

operations, functions, conditions, states, events, workflows and other elements (collectively encompassed by the term "states") of such environments. Data may also include metadata about any of the foregoing, such as describing data, indicating provenance, indicating elements relating to identity, access, roles, and permissions, providing summaries or abstractions of data, or otherwise augmenting one or more items of data to enable further processing, such as for extraction, transforming, loading, and processing data. Such data (such term including metadata except where context indicates otherwise) may be highly valuable to third parties, either as an individual element (such as the instance where data about the state of an environment can be used as a condition within a process) or in the aggregate (such as the instance where collected data, optionally over many systems and devices in different environments can be used to develop models of behavior, to train learning systems, or the like). As billions of IoT devices are deployed, with countless connections, the amount of available data will proliferate. To enable access and utilization of data, the cognitive data marketplace 4102 enables various components, features, services, and processes for enabling users to supply, find, consume, and transact in packages of data, such as batches of data, streams of data (including event streams), data from various data pools 4120, and the like. In example embodiments, the cognitive data marketplace 4102 may be included in, connected to, or integrated with, one or more other components of a host processing architecture 4024 of a host processing system 112, such as a cloud-based system, as well as to various sensors, input sources 115, data collection systems 102 and the like. The cognitive data marketplace 4102 may include marketplace interfaces 4108, which may include one or more supplier interfaces by which data suppliers may make data available and one more consumer interfaces by which data may be found and acquired. The consumer interface may include an interface to a data market search system 4118, which may include features that enable a user to indicate what types of data a user wishes to obtain, such as by entering keywords in a natural language search interface that characterize data or metadata. The search interface can use various search and filtering techniques, including keyword matching, collaborative filtering (such as using known preferences or characteristics of the consumer to match to similar consumers and the past outcomes of those other consumers), ranking techniques (such as ranking based on success of past outcomes according to various metrics, such as those described in connection with other embodiments in this disclosure). In example embodiments, a supply interface may allow an owner or supplier of data to supply the data in one or more packages to and through the cognitive data marketplace 4102, such as packaging batches of data, streams of data, or the like. The supplier may pre-package data, such as by providing data from a single input source 116, a single sensor, and the like, or by providing combinations, permutations, and the like (such as multiplexed analog data, mixed bytes of data from multiple sources, results of extraction, loading and transformation, results of convolution, and the like), as well as by providing metadata with respect to any of the foregoing. Packaging may include pricing, such as on a per-batch basis, on a streaming basis (such as subscription to an event feed or other feed or stream), on a per item basis, on a revenue share basis, or other basis. For data involving pricing, a data transaction system 4114 may track orders, delivery, and utilization, including fulfillment of orders. The transaction system 4114 may include

rich transaction features, including digital rights management, such as by managing cryptographic keys that govern access control to purchased data, that govern usage (such as allowing data to be used for a limited time, in a limited domain, by a limited set of users or roles, or for a limited purpose). The transaction system 4114 may manage payments, such as by processing credit cards, wire transfers, debits, and other forms of consideration.

[0358] In example embodiments, a cognitive data packaging system 4012 of the marketplace 4102 may use machine-based intelligence to package data, such as by automatically configuring packages of data in batches, streams, pools, or the like. In example embodiments, packaging may be according to one or more rules, models, or parameters, such as by packaging or aggregating data that is likely to supplement or complement an existing model. For example, operating data from a group of similar machines (such as one or more industrial machines noted throughout this disclosure) may be aggregated together, such as based on metadata indicating the type of data or by recognizing features or characteristics in the data stream that indicate the nature of the data. In example embodiments, packaging may occur using machine learning and cognitive capabilities, such as by learning what combinations, permutations, mixes, layers, and the like of input sources 116, sensors, information from data pools 4120 and information from data collection systems 102 are likely to satisfy user requirements or result in measures of success. Learning may be based on learning feedback 4012, such as learning based on measures determined in an analytic system 4018, such as system performance measures, data collection measures, analytic measures, and the like. In example embodiments, success measures may be correlated to marketplace success measures, such as viewing of packages, engagement with packages, purchase or licensing of packages, payments made for packages, and the like. Such measures may be calculated in an analytic system 4018, including associating particular feedback measures with search terms and other inputs, so that the cognitive packaging system 4110 can find and configure packages that are designed to provide increased value to consumers and increased returns for data suppliers. In example embodiments, the cognitive data packaging system 4110 can automatically vary packaging, such as using different combinations, permutations, mixes, and the like, and varying weights applied to given input sources, sensors, data pools and the like, using learning feedback 4012 to promote favorable packages and de-emphasize less favorable packages. This may occur using genetic programming and similar techniques that compare outcomes for different packages. Feedback may include state information from the state system 4020 (such as about various operating states, and the like), as well as about marketplace conditions and states, such as pricing and availability information for other data sources. Thus, an adaptive cognitive data packaging system 4110 is provided that automatically adapts to conditions to provide favorable packages of data for the marketplace 4102.

[0359] In example embodiments, a cognitive data pricing system 4112 may be provided to set pricing for data packages. In example embodiments, the data pricing system 4112 may use a set of rules, models, or the like, such as setting pricing based on supply conditions, demand conditions, pricing of various available sources, and the like. For example, pricing for a package may be configured to be set based on the sum of the prices of constituent elements (such as input sources,

sensor data, or the like), or to be set based on a rule-based discount to the sum of prices for constituent elements, or the like. Rules and conditional logic may be applied, such as rules that factor in cost factors (such as bandwidth and network usage, peak demand factors, scarcity factors, and the like), rules that factor in utilization parameters (such as the purpose, domain, user, role, duration, or the like for a package) and many others. In example embodiments, the cognitive data pricing system 4112 may include fully cognitive, intelligent features, such as using genetic programming including automatically varying pricing and tracking feedback on outcomes. Outcomes on which tracking feedback may be based include various financial yield metrics, utilization metrics and the like that may be provided by calculating metrics in an analytic system 4018 on data from the data transaction system 4114.

[0360] Methods and systems are disclosed herein for self-organizing data pools which may include self-organization of data pools based on utilization and/or yield metrics, including utilization and/or yield metrics that are tracked for a plurality of data pools. The data pools may initially comprise unstructured or loosely structured pools of data that contain data from industrial environments, such as sensor data from or about industrial machines or components. For example, a data pool might take streams of data from various machines or components in an environment, such as turbines, compressors, batteries, reactors, engines, motors, vehicles, pumps, rotors, axles, bearings, valves, and many others, with the data streams containing analog and/or digital sensor data (of a wide range of types), data published about operating conditions, diagnostic and fault data, identifying data for machines or components, asset tracking data, and many other types of data. Each stream may have an identifier in the pool, such as indicating its source, and optionally its type. The data pool may be accessed by external systems, such as through one or more interfaces or APIs (e.g., RESTful APIs), or by data integration elements (such as gateways, brokers, bridges, connectors, or the like), and the data pool may use similar capabilities to get access to available data streams. A data pool may be managed by a self-organizing machine learning facility, which may configure the data pool, such as by managing what sources are used for the pool, managing what streams are available, and managing APIs or other connections into and out of the data pool. The self-organization may take feedback such as based on measures of success that may include measures of utilization and yield. The measures of utilization and yield that may include may account for the cost of acquiring and/or storing data, as well as the benefits of the pool, measured either by profit or by other measures that may include user indications of usefulness, and the like. For example, a self-organizing data pool might recognize that chemical and radiation data for an energy production environment are regularly accessed and extracted, while vibration and temperature data have not been used, in which case the data pool might automatically reorganize, such as by ceasing storage of vibration and/or temperature data, or by obtaining better sources of such data. This automated reorganization can also apply to data structures, such as promoting different data types, different data sources, different data structures, and the like, through progressive iteration and feedback.

[0361] In example embodiments, a platform is provided having self-organization of data pools based on utilization and/or yield metrics. In example embodiments, the data pools 4020 may be

self-organizing data pools 4020, such as being organized by cognitive capabilities as described throughout this disclosure. The data pools 4020 may self-organize in response to learning feedback 4012, such as based on feedback of measures and results, including calculated in an analytic system 4018. Organization may include determining what data or packages of data to store in a pool (such as representing particular combinations, permutations, aggregations, and the like), the structure of such data (such as in flat, hierarchical, linked, or other structures), the duration of storage, the nature of storage media (such as hard disks, flash memory, SSDs, network-based storage, or the like), the arrangement of storage bits, and other parameters. The content and nature of storage may be varied, such that a data pool 4020 may learn and adapt, such as based on states of the host system 112, one or more data collection systems 102, storage environment parameters (such as capacity, cost, and performance factors), data collection environment parameters, marketplace parameters, and many others. In example embodiments, pools 4020 may learn and adapt, such as by variation of the above and other parameters in response to yield metrics (such as return on investment, optimization of power utilization, optimization of revenue, and the like).

[0362] Methods and systems are disclosed herein for training AI models based on industry-specific feedback, including training an AI model based on industry-specific feedback that reflects a measure of utilization, yield, or impact, and where the AI model operates on sensor data from an industrial environment. As noted above, these models may include operating models for industrial environments, machines, workflows, models for anticipating states, models for predicting fault and optimizing maintenance, models for self-organizing storage (on devices, in data pools and/or in the cloud), models for optimizing data transport (such as for optimizing network coding, network-condition-sensitive routing, and the like), models for optimizing data marketplaces, and many others.

[0363] In example embodiments, a platform is provided having training AI models based on industry-specific feedback. In example embodiments, the various embodiments of cognitive systems disclosed herein may take inputs and feedback from industry-specific and domain-specific sources 116 (such as relating to optimization of specific machines, devices, components, processes, and the like). Thus, learning and adaptation of storage organization, network usage, combination of sensor and input data, data pooling, data packaging, data pricing, and other features (such as for a marketplace 4102 or for other purposes of the host processing system 112) may be configured by learning on the domain-specific feedback measures of a given environment or application, such as an application involving IoT devices (such as an industrial environment). This may include optimization of efficiency (such as in electrical, electromechanical, magnetic, physical, thermodynamic, chemical and other processes and systems), optimization of outputs (such as for production of energy, materials, products, services and other outputs), prediction, avoidance and mitigation of faults (such as in the aforementioned systems and processes), optimization of performance measures (such as returns on investment, yields, profits, margins, revenues and the like), reduction of costs (including labor costs, bandwidth costs, data costs, material input costs, licensing costs, and many others), optimization of benefits (such as relating to safety, satisfaction,

health), optimization of work flows (such as optimizing time and resource allocation to processes), and others.

[0364] Methods and systems are disclosed herein for a self-organized swarm of industrial data collectors, including a self-organizing swarm of industrial data collectors that organize among themselves to optimize data collection based on the capabilities and conditions of the members of the swarm. Each member of the swarm may be configured with intelligence, and the ability to coordinate with other members. For example, a member of the swarm may track information about what data other members are handling, so that data collection activities, data storage, data processing, and data publishing can be allocated intelligently across the swarm, taking into account conditions of the environment, capabilities of the members of the swarm, operating parameters, rules (such as from a rules engine that governs the operation of the swarm), and current conditions of the members. For example, among four collectors, one that has relatively low current power levels (such as a low battery), might be temporarily allocated the role of publishing data, because it may receive a dose of power from a reader or interrogation device (such as an RFID reader) when it needs to publish the data. A second collector with good power levels and robust processing capability might be assigned more complex functions, such as processing data, fusing data, organizing the rest of the swarm (including self-organization under machine learning, such that the swarm is optimized over time, including by adjusting operating parameters, rules, and the like based on feedback), and the like. A third collector in the swarm with robust storage capabilities might be assigned the task of collecting and storing a category of data, such as vibration sensor data, that consumes considerable bandwidth. A fourth collector in the swarm, such as one with lower storage capabilities, might be assigned the role of collecting data that can usually be discarded, such as data on current diagnostic conditions, where only data on faults needs to be maintained and passed along. Members of a swarm may connect by peer-to-peer relationships by using a member as a "master" or "hub," or by having them connect in a series or ring, where each member passes along data (including commands) to the next, and is aware of the nature of the capabilities and commands that are suitable for the preceding and/or next member. The swarm may be used for allocation of storage across it (such as using memory of each memory as an aggregate data store. In these examples, the aggregate data store may support a distributed ledger, which may store transaction data, such as for transactions involving data collected by the swarm, transactions occurring in the industrial environment, or the like. In example embodiments, the transaction data may also include data used to manage the swarm, the environment, or a machine or components thereof. The swarm may self-organize, either by machine learning capability disposed on one or more members of the swarm, or based on instructions from an external machine learning facility, which may optimize storage, data collection, data processing, data presentation, data transport, and other functions based on managing parameters that are relevant to each. The machine learning facility may start with an initial configuration and vary parameters of the swarm relevant to any of the foregoing (also including varying the membership of the swarm), such as iterating based on feedback to the machine learning facility regarding measures of success (such as utilization measures, efficiency measures, measures of success in prediction or anticipation of states,

productivity measures, yield measures, profit measures, and others). Over time, the swarm may be optimized to a favorable configuration to achieve the desired measure of success for an owner, operator, or host of an industrial environment or a machine, component, or process thereof.

[0365] The swarm 4202 may be organized based on a hierarchical organization (such as where a master data collector 102 organizes and directs activities of one or more subservient data collectors 102), a collaborative organization (such as where decision-making for the organization of the swarm 4202 is distributed among the data collectors 102 (such as using various models for decision-making, such as voting systems, points systems, least-cost routing systems, prioritization systems, and the like), and the like.) In example embodiments, one or more of the data collectors 102 may have mobility capabilities, such as in cases where a data collector is disposed on or in a mobile robot, drone, mobile submersible, or the like, so that organization may include the location and positioning of the data collectors 102. Data collection systems 102 may communicate with each other and with the host processing system 112, including sharing an aggregate allocated storage space involving storage on or accessible to one or more of the collectors (which in embodiment may be treated as a unified storage space even if physically distributed, such as using virtualization capabilities). Organization may be automated based on one or more rules, models, conditions, processes, or the like (such as embodied or executed by conditional logic), and organization may be governed by policies, such as handled by the policy engine. Rules may be based on industry, application- and domain-specific objects, classes, events, workflows, processes, and systems, such as by setting up the swarm 4202 to collect selected types of data at designated places and times, such as coordinated with the foregoing. For example, the swarm 4202 may assign data collectors 102 to serially collect diagnostic, sensor, instrumentation and/or telematic data from each of a series of machines that execute an industrial process (such as a robotic manufacturing process), such as at the time and location of the input to and output from each of those machines. In example embodiments, self-organization may be cognitive, such as where the swarm varies one or more collection parameters and adapts the selection of parameters, weights applied to the parameters, or the like, over time. In examples, this may be in response to learning and feedback, such as from the learning feedback system 4012 that may be based on various feedback measures that may be determined by applying the analytic system 4018 (which in embodiments may reside on the swarm 4202, the host processing system 112, or a combination thereof) to data handled by the swarm 4202 or to other elements of the various embodiments disclosed herein (including marketplace elements and others). Thus, the swarm 4202 may display adaptive behavior, such as adapting to the current state 4020 or an anticipated state of its environment (accounting for marketplace behavior), behavior of various objects (such as IoT devices, machines, components, and systems), processes (including events, states, workflows, and the like), and other factors at a given time. Parameters that may be varied in a process of variation (such as in a neural net, self-organizing map, or the like), selection, promotion, or the like (such as those enabled by genetic programming or other AI-based techniques). Parameters that may be managed, varied, selected and adapted by cognitive, machine learning may include storage parameters (location, type, duration, amount, structure and the like across the swarm 4202), network parameters (such as how

the swarm 4202 is organized, such as in mesh, peer-to-peer, ring, serial, hierarchical and other network configurations as well as bandwidth utilization, data routing, network protocol selection, network coding type, and other networking parameters), security parameters (such as settings for various security applications and services), location and positioning parameters (such as routing movement of mobile data collectors 102 to locations, positioning and orienting collectors 102 and the like relative to points of data acquisition, relative to each other, and relative to locations where network availability may be favorable, among others), input selection parameters (such as input selection among sensors, input sources 116 and the like for each collector 102 and for the aggregate collection), data combination parameters (such as those for sensor fusion, input combination, multiplexing, mixing, layering, convolution, and other combinations), power parameters (such as parameters based on power levels and power availability for one or more collectors 102 or other objects, devices, or the like), states (including anticipated states and conditions of the swarm 4202, individual collection systems 102, the host processing system 112 or one or more objects in an environment), events, and many others. Feedback may be based on any of the kinds of feedback described herein, such that over time the swarm may adapt to its current and anticipated situation to achieve a wide range of desired objectives.

[0366] Methods and systems are disclosed herein for an industrial IoT distributed ledger, including a distributed ledger supporting the tracking of transactions executed in an automated data marketplace for industrial IoT data. A distributed ledger may distribute storage across devices, using a secure protocol, such as those used for cryptocurrencies (such as the Blockchain™ protocol used to support the Bitcoin™ currency). A ledger or similar transaction record, which may comprise a structure where each successive member of a chain stores data for previous transactions, and a competition can be established to determine which of alternative data stored data structures is "best" (such as being most complete), can be stored across data collectors, industrial machines or components, data pools, data marketplaces, cloud computing elements, servers, and/or on the IT infrastructure of an enterprise (such as an owner, operator or host of an industrial environment or of the systems disclosed herein). The ledger or transaction may be optimized by machine learning, such as to provide storage efficiency, security, redundancy, or the like.

[0367] In example embodiments, the cognitive data marketplace 4102 may use a secure architecture for tracking and resolving transactions, such as a distributed ledger 4004, wherein transactions in data packages are tracked in a chained, distributed data structure, such as a Blockchain™, allowing forensic analysis and validation where individual devices store a portion of the ledger representing transactions in data packages. The distributed ledger 4004 may be distributed to IoT devices, to data pools 4020, to data collection systems 102, and the like, so that transaction information can be verified without reliance on a single, central repository of information. The transaction system 4114 may be configured to store data in the distributed ledger 4004 and to retrieve data from it (and from constituent devices) in order to resolve transactions. Thus, a distributed ledger 4004 for handling transactions in data, such as for packages of IoT data, is provided. In example embodiments, the self-organizing storage system 4028 may be used for

optimizing storage of distributed ledger data, as well as for organizing storage of packages of data, such as IoT data, that can be presented in the marketplace 4102.

[0368] Methods and systems are disclosed herein for a network-sensitive collector, including a network condition-sensitive, self-organizing, multi-sensor data collector that can optimize based on bandwidth, quality of service, pricing and/or other network conditions. Network sensitivity can include awareness of the price of data transport (such as allowing the system to pull or push data during off-peak periods or within the available parameters of paid data plans), the quality of the network (such as to avoid periods where errors are likely), the quality of environmental conditions (such as delaying transmission until signal quality is good, such as when a collector emerges from a shielded environment, avoiding wasting use of power when seeking a signal when shielded, such as by large metal structures typically of industrial environments), and the like.

[0369] Methods and systems are disclosed herein for a remotely organized universal data collector that can power up and down sensor interfaces based on need and/or conditions identified in an industrial data collection environment. For example, interfaces can recognize what sensors are available and interfaces and/or processors can be turned on to take input from such sensors, including hardware interfaces that allow the sensors to plug in to the data collector, wireless data interfaces (such as where the collector can ping the sensor, optionally providing some power via an interrogation signal), and software interfaces (such as for handling particular types of data). Thus, a collector that is capable of handling various kinds of data can be configured to adapt to the particular use in a given environment. In example embodiments, configuration may be automatic or under machine learning, which may improve configuration by optimizing parameters based on feedback measures over time.

[0370] Methods and systems are disclosed herein for self-organizing storage for a multi-sensor data collector, including self-organizing storage for a multi-sensor data collector for industrial sensor data. Self-organizing storage may allocate storage based on application of machine learning, which may improve storage configuration based on feedback measure over time. Storage may be optimized by configuring what data types are used (e.g., byte-like structures, structures representing fused data from multiple sensors, structures representing statistics or measures calculated by applying mathematical functions on data, and the like), by configuring compression, by configuring data storage duration, by configuring write strategies (such as by striping data across multiple storage devices, using protocols where one device stores instructions for other devices in a chain, and the like), and by configuring storage hierarchies, such as by providing pre-calculated intermediate statistics to facilitate more rapid access to frequently accessed data items. Thus, highly intelligent storage systems may be configured and optimized, based on feedback, over time.

[0371] Methods and systems are disclosed herein for self-organizing network coding for a multi-sensor data network, including self-organizing network coding for a data network that transports data from multiple sensors in an industrial data collection environment. Network coding, including random linear network coding, can enable highly efficient and reliable transport of large amounts of data over various kinds of networks. Different network coding configurations can be selected,

based on machine learning, to optimize network coding and other network transport characteristics based on network conditions, environmental conditions, and other factors, such as the nature of the data being transported, environmental conditions, operating conditions, and the like (including by training a network coding selection model over time based on feedback of measures of success, such as any of the measures described herein).

[0372] In example embodiments, a platform is provided having a self-organizing network coding for multi-sensor data network. A cognitive system may vary one or more parameters for networking, such as network type selection (e.g., selecting among available local, cellular, satellite, Wi-Fi, BLUETOOTH™, NFC, ZIGBEE® and other networks), network selection (such as selecting a specific network, such as one that is known to have desired security features), network coding selection (such as selecting a type of network coding for efficient transport[such as random linear network coding, fixed coding, and others]), network timing selection (such as configuring delivery based on network pricing conditions, traffic and the like), network feature selection (such as selecting cognitive features, security features, and the like), network conditions (such as network quality based on current environmental or operation conditions), network feature selection (such as enabling available authentication, permission and similar systems), network protocol selection (such as among HTTP, IP, TCP/IP, cellular, satellite, serial, packet, streaming, and many other protocols), and others. Given bandwidth constraints, price variations, sensitivity to environmental factors, security concerns, and the like, selecting the optimal network configuration can be highly complex and situation dependent. The self-organizing networking system 4030 may vary combinations and permutations of these parameters while taking input from a learning feedback system 4012 such as using information from the analytic system 4018 about various measures of outcomes. In the many examples, outcomes may include overall system measures, analytic success measures, and local performance indicators. In example embodiments, input from a learning feedback system 4012 may include information from various sensors and input sources 116, information from the state system 4020 about states (such as events, environmental conditions, operating conditions, and many others, or other information) or taking other inputs. By variation and selection of alternative configurations of networking parameters in different states, the self-organizing networking system may find configurations that are well-adapted to the environment that is being monitored or controlled by the host system 112, such as the instance where one or more data collection systems 102 are located and that are well-adapted to emerging network conditions. Thus, a self-organizing, network-condition-adaptive data collection system is provided.

[0373] Referring to Figure 32, a data collection system 102 may have one or more output interfaces and/or ports 4010. These may include network ports and connections, application programming interfaces, and the like. Methods and systems are disclosed herein for a haptic or multi-sensory user interface, including a wearable haptic or multi-sensory user interface for an industrial sensor data collector, with vibration, heat, electrical, and/or sound outputs. For example, an interface may, based on a data structure configured to support the interface, be set up to provide a user with input or feedback, such as based on data from sensors in the environment. For example, if a fault condition based on a vibration data (such as resulting from a bearing being worn down, an axle

being misaligned, or a resonance condition between machines) is detected, it can be presented in a haptic interface by vibration of an interface, such as shaking a wrist-worn device. Similarly, thermal data indicating overheating could be presented by warming or cooling a wearable device, such as while a worker is working on a machine and cannot necessarily look at a user interface. Similarly, electrical or magnetic data may be presented by a buzzing, and the like, such as to indicate presence of an open electrical connection or wire, etc. That is, a multi-sensory interface can intuitively help a user (such as a user with a wearable device) get a quick indication of what is going on in an environment, with the wearable interface having various modes of interaction that do not require a user to have eyes on a graphical UI, which may be difficult or impossible in many industrial environments where a user needs to keep an eye on the environment.

[0374] In example embodiments, a platform is provided having a wearable haptic user interface for an industrial sensor data collector, with vibration, heat, electrical, and/or sound outputs. In example embodiments, a haptic user interface 4302 is provided as an output for a data collection system 102, such as a system for handling and providing information for vibration, heat, electrical, and/or sound outputs, such as to one or more components of the data collection system 102 or to another system, such as a wearable device, mobile phone, or the like. A data collection system 102 may be provided in a form factor suitable for delivering haptic input to a user, such as vibration, warming or cooling, buzzing, or the like, such as input disposed in headgear, an armband, a wristband or watch, a belt, an item of clothing, a uniform, or the like. In such cases, data collection systems 102 may be integrated with gear, uniforms, equipment, or the like worn by users, such as individuals responsible for operating or monitoring an industrial environment. In example embodiments, signals from various sensors or input sources (or selective combinations, permutations, mixes, and the like, as managed by one or more of the cognitive input selection systems 4004, 4014) may trigger haptic feedback. For example, if a nearby industrial machine is overheating, the haptic interface may alert a user by warming up, or by sending a signal to another device (such as a mobile phone) to warm up. If a system is experiencing unusual vibrations, the haptic interface may vibrate. Thus, through various forms of haptic input, a data collection system 102 may inform users of the need to attend to one or more devices, machines, or other factors (such as those in an industrial environment) without requiring them to read messages or divert their visual attention away from the task at hand. The haptic interface, and selection of what outputs should be provided, may be considered in the cognitive input selection systems 4004, 4014. For example, user behavior (such as responses to inputs) may be monitored and analyzed in an analytic system 4018, and feedback may be provided through the learning feedback system 4012, so that signals may be provided based on the right collection or package of sensors and inputs, at the right time and in the right manner, to optimize the effectiveness of the haptic system 4202. This may include rule-based or model-based feedback (such as providing outputs that correspond in some logical fashion to the source data that is being conveyed). In example embodiments, a cognitive haptic system may be provided, where selection of inputs or triggers for haptic feedback, selection of outputs, timing, intensity levels, durations, and other parameters (or weights applied to them) may be varied in a process of variation, promotion, and selection (such as using genetic programming)

with feedback based on real world responses to feedback in actual situations or based on results of simulation and testing of user behavior. Thus, an adaptive haptic interface for a data collection system 102 is provided, which may learn and adapt feedback to satisfy requirements and to optimize the impact on user behavior, such as for overall system outcomes, data collection outcomes, analytic outcomes, and the like.

[0375] Methods and systems are disclosed herein for a presentation layer for AR/VR industrial glasses, where heat map elements are presented based on patterns and/or parameters in collected data. Methods and systems are disclosed herein for condition-sensitive, self-organized tuning of AR/VR interfaces based on feedback metrics and/or training in industrial environments. In example embodiments, any of the data, measures, and the like described throughout this disclosure can be presented by visual elements, overlays, and the like for presentation in the AR/VR interfaces, such as in industrial glasses, on AR/VR interfaces on smart phones or tablets, on AR/VR interfaces on data collectors (which may be embodied in smart phones or tablets), on displays located on machines or components, and/or on displays located in industrial environments.

[0376] In example embodiments, a platform is provided having heat maps displaying collected data for AR/VR. In example embodiments, a platform is provided having heat maps 4204 displaying collected data from a data collection system 102 for providing input to an AR/VR interface 4208. In example embodiments, the heat map interface 4304 is provided as an output for a data collection system 102, such as for handling and providing information for visualization of various sensor data and other data (such as map data, analog sensor data, and other data), such as to one or more components of the data collection system 102 or to another system, such as a mobile device, tablet, dashboard, computer, AR/VR device, or the like. A data collection system 102 may be provided in a form factor suitable for delivering visual input to a user, such as the presentation of a map that includes indicators of levels of analog and digital sensor data (such as data indicating levels of rotation, vibration, heating or cooling, pressure, and many other conditions). In such cases, data collection systems 102 may be integrated with equipment, or the like, that are used by individuals responsible for operating or monitoring an industrial environment. In example embodiments, signals from various sensors or input sources (or selective combinations, permutations, mixes, and the like, as managed by one or more of the cognitive input selection systems 4004, 4014) may provide input data to a heat map. Coordinates may include real world location coordinates (such as geo-location or location on a map of an environment), as well as other coordinates, such as time-based coordinates, frequency-based coordinates, or other coordinates that allow for representation of analog sensor signals, digital signals, input source information, and various combinations, in a map-based visualization, such that colors may represent varying levels of input along the relevant dimensions. For example, if a nearby industrial machine is overheating, the heat map interface may alert a user by showing a machine in bright red. If a system is experiencing unusual vibrations, the heat map interface may show a different color for a visual element for the machine, or it may cause an icon or display element representing the machine to vibrate in the interface, calling attention to the element. Clicking, touching, or otherwise interacting with the map can allow a user to drill down and see underlying sensor or input data that is used as

an input to the heat map display. Thus, through various forms of display, a data collection system 102 may inform users of the need to attend to one or more devices, machines, or other factors, such as those in an industrial environment, without requiring them to read text-based messages or input. The heat map interface, and selection of what outputs should be provided, may be considered in the cognitive input selection systems 4004, 4014. For example, user behavior (such as responses to inputs or displays) may be monitored and analyzed in an analytic system 4018, and feedback may be provided through the learning feedback system 4012, so that signals may be provided based on the right collection or package of sensors and inputs, at the right time and in the right manner, to optimize the effectiveness of the heat map UI 4304. This may include rule-based or model-based feedback (such as feedback providing outputs that correspond in some logical fashion to the source data that is being conveyed). In example embodiments, a cognitive heat map system may be provided, where selection of inputs or triggers for heat map displays, selection of outputs, colors, visual representation elements, timing, intensity levels, durations and other parameters (or weights applied to them) may be varied in a process of variation, promotion and selection (such as selection using genetic programming) with feedback based on real world responses to feedback in actual situations or based on results of simulation and testing of user behavior. Thus, an adaptive heat map interface for a data collection system 102, or data collected thereby 102, or data handled by a host processing system 112, is provided, which may learn and adapt feedback to satisfy requirements and to optimize the impact on user behavior and reaction, such as for overall system outcomes, data collection outcomes, analytic outcomes, and the like.

[0377] In example embodiments, a platform is provided having automatically tuned AR/VR visualization of data collected by a data collector. In example embodiments, a platform is provided having an automatically tuned AR/VR visualization system 4308 for visualization of data collected by a data collection system 102, such as the case where the data collection system 102 has an AR/VR interface 4208 or provides input to an AR/VR interface 4308 (such as a mobile phone positioned in a virtual reality or AR headset, a set of AR glasses, or the like). In example embodiments, the AR/VR system 4308 is provided as an output interface of a data collection system 102, such as a system for handling and providing information for visualization of various sensor data and other data (such as map data, analog sensor data, and other data), such as to one or more components of the data collection system 102 or to another system, such as a mobile device, tablet, dashboard, computer, AR/VR device, or the like. A data collection system 102 may be provided in a form factor suitable for delivering AR or VR visual, auditory, or other sensory input to a user, such as by presenting one or more displays such as 3D-realistic visualizations, objects, maps, camera overlays, or other overlay elements, maps and the like that include or correspond to indicators of levels of analog and digital sensor data (such as data indicating levels of rotation, vibration, heating or cooling, pressure and many other conditions, to input sources 116, or the like). In such cases, data collection systems 102 may be integrated with equipment, or the like, that are used by individuals responsible for operating or monitoring an industrial environment.

[0378] In example embodiments, signals from various sensors or input sources (or selective combinations, permutations, mixes, and the like as managed by one or more of the cognitive input

selection systems 4004, 4014) may provide input data to populate, configure, modify, or otherwise determine the AR/VR element. Visual elements may include a wide range of icons, map elements, menu elements, sliders, toggles, colors, shapes, sizes, and the like, for representation of analog sensor signals, digital signals, input source information, and various combinations. In many examples, colors, shapes, and sizes of visual overlay elements may represent varying levels of input along the relevant dimensions for a sensor or combination of sensors. In further examples, if a nearby industrial machine is overheating, an AR element may alert a user by showing an icon representing that type of machine in flashing red color in a portion of the display of a pair of AR glasses. If a system is experiencing unusual vibrations, a virtual reality interface showing visualization of the components of the machine (such as an overlay of a camera view of the machine with 3D visualization elements) may show a vibrating component in a highlighted color, with motion, or the like, to ensure the component stands out in a virtual reality environment being used to help a user monitor or service the machine. Clicking, touching, moving eyes toward, or otherwise interacting with a visual element in an AR/VR interface may allow a user to drilldown and see underlying sensor or input data that is used as an input to the display. Thus, through various forms of display, a data collection system 102 may inform users of the need to attend to one or more devices, machines, or other factors (such as in an industrial environment), without requiring them to read text-based messages or input or divert attention from the applicable environment (whether it is a real environment with AR features or a virtual environment, such as for simulation, training, or the like).

[0379] The AR/VR output interface 4208, and selection and configuration of what outputs or displays should be provided, may be handled in the cognitive input selection systems 4004, 4014. For example, user behavior (such as responses to inputs or displays) may be monitored and analyzed in an analytic system 4018, and feedback may be provided through the learning feedback system 4012, so that AR/VR display signals may be provided based on the right collection or package of sensors and inputs, at the right time and in the right manner, to optimize the effectiveness of the AR/VR UI 4308. This may include rule-based or model-based feedback (such as providing outputs that correspond in some logical fashion to the source data that is being conveyed). In example embodiments, a cognitively tuned AR/VR interface control system 4308 may be provided, where selection of inputs or triggers for AR/VR display elements, selection of outputs (such as colors, visual representation elements, timing, intensity levels, durations and other parameters [or weights applied to them]) and other parameters of a VR/AR environment may be varied in a process of variation, promotion and selection (such as the use of genetic programming) with feedback based on real world responses in actual situations or based on results of simulation and testing of user behavior. Thus, an adaptive, tuned AR/VR interface for a data collection system 102, or data collected thereby 102, or data handled by a host processing system 112, is provided, which may learn and adapt feedback to satisfy requirements and to optimize the impact on user behavior and reaction, such as for overall system outcomes, data collection outcomes, analytic outcomes, and the like.

[0380] As noted above, methods and systems are disclosed herein for continuous ultrasonic monitoring, including providing continuous ultrasonic monitoring of rotating elements and bearings of an energy production facility. Embodiments include using continuous ultrasonic monitoring of an industrial environment as a source for a cloud-deployed pattern recognizer. Embodiments include using continuous ultrasonic monitoring to provide updated state information to a state machine that is used as an input to a cloud-deployed pattern recognizer. Embodiments include making available continuous ultrasonic monitoring information to a user based on a policy declared in a policy engine. Embodiments include storing continuous ultrasonic monitoring data with other data in a fused data structure on an industrial sensor device. Embodiments include making a stream of continuous ultrasonic monitoring data from an industrial environment available as a service from a data marketplace. Embodiments include feeding a stream of continuous ultrasonic monitoring data into a self-organizing data pool. Embodiments include training a machine learning model to monitor a continuous ultrasonic monitoring data stream where the model is based on a training set created from human analysis of such a data stream, and is improved based on data collected on performance in an industrial environment.

[0381] Embodiments include a swarm of data collectors that include at least one data collector for continuous ultrasonic monitoring of an industrial environment and at least one other type of data collector. Embodiments include using a distributed ledger to store time-series data from continuous ultrasonic monitoring across multiple devices. Embodiments include collecting a stream of continuous ultrasonic data in a self-organizing data collector, a network-sensitive data collector, a remotely organized data collector, a data collector having self-organized storage and the like. Embodiments include using self-organizing network coding to transport a stream of ultrasonic data collected from an industrial environment. Embodiments include conveying an indicator of a parameter of a continuously collected ultrasonic data stream via an interface where the interface is one of a sensory interface of a wearable device, a heat map visual interface of a wearable device, an interface that operates with self-organized tuning of the interface layer, and the like.

[0382] As noted above, methods and systems are disclosed herein for cloud-based, machine pattern recognition based on fusion of remote analog industrial sensors. Embodiments include taking input from a plurality of analog sensors disposed in an industrial environment, multiplexing the sensors into a multiplexed data stream, feeding the data stream into a cloud-deployed machine learning facility, and training a model of the machine learning facility to recognize a defined pattern associated with the industrial environment. Embodiments include using a cloud-based pattern recognizer on input states from a state machine that characterizes states of an industrial environment. Embodiments include deploying policies by a policy engine that govern what data can be used by what users and for what purpose in cloud-based, machine learning. Embodiments include using a cloud-based platform to identify patterns in data across a plurality of data pools that contain data published from industrial sensors. Embodiments include training a model to identify preferred sensor sets to diagnose a condition of an industrial environment, where a training set is created by a human user and the model is improved based on feedback from data collected about conditions in an industrial environment.

[0383] Embodiments include a swarm of data collectors that is governed by a policy that is automatically propagated through the swarm. Embodiments include using a distributed ledger to store sensor fusion information across multiple devices. Embodiments include feeding input from a set of data collectors into a cloud-based pattern recognizer that uses data from multiple sensors for an industrial environment. The data collectors may be self-organizing data collectors, network-sensitive data collectors, remotely organized data collectors, a set of data collectors having self-organized storage, and the like. Embodiments include a system for data collection in an industrial environment with self-organizing network coding for data transport of data fused from multiple sensors in the environment. Embodiments include conveying information formed by fusing inputs from multiple sensors in an industrial data collection system in an interface such as a multi-sensory interface, a heat map interface, an interface that operates with self-organized tuning of the interface layer, and the like.

[0384] As noted above, methods and systems are disclosed herein for cloud-based, machine pattern analysis of state information from multiple analog industrial sensors to provide anticipated state information for an industrial system. Embodiments include using a policy engine to determine what state information can be used for cloud-based machine analysis. Embodiments include feeding inputs from multiple devices that have fused and on-device storage of multiple sensor streams into a cloud-based pattern recognizer to determine an anticipated state of an industrial environment. Embodiments include making an output, such as anticipated state information, from a cloud-based machine pattern recognizer that analyzes fused data from remote, analog industrial sensors available as a data service in a data marketplace. Embodiments include using a cloud-based pattern recognizer to determine an anticipated state of an industrial environment based on data collected from data pools that contain streams of information from machines in the environment. Embodiments include training a model to identify preferred state information to diagnose a condition of an industrial environment, where a training set is created by a human user and the model is improved based on feedback from data collected about conditions in an industrial environment. Embodiments include a swarm of data collectors that feeds a state machine that maintains current state information for an industrial environment. Embodiments include using a distributed ledger to store historical state information for fused sensor states a self-organizing data collector that feeds a state machine that maintains current state information for an industrial environment. Embodiments include a data collector that feeds a state machine that maintains current state information for an industrial environment where the data collector may be a network sensitive data collector, a remotely organized data collector, a data collector with self-organized storage, and the like. Embodiments include a system for data collection in an industrial environment with self-organizing network coding for data transport and maintains anticipated state information for the environment. Embodiments include conveying anticipated state information determined by machine learning in an industrial data collection system in an interface where the interface may be one or more of a multisensory interface, a heat map interface an interface that operates with self-organized tuning of the interface layer, and the like.

[0385] As noted above, methods and systems are disclosed herein for a cloud-based policy automation engine for IoT, with creation, deployment, and management of IoT devices, including a cloud-based policy automation engine for IoT, enabling creation, deployment and management of policies that apply to IoT devices. Policies can relate to data usage to an on-device storage system that stores fused data from multiple industrial sensors, or what data can be provided to whom in a self-organizing marketplace for IoT sensor data. Policies can govern how a self-organizing swarm or data collector should be organized for a particular industrial environment, how a network-sensitive data collector should use network bandwidth for a particular industrial environment, how a remotely organized data collector should collect, and make available, data relating to a specified industrial environment, or how a data collector should self-organize storage for a particular industrial environment. Policies can be deployed across a set of self-organizing pools of data that contain data streamed from industrial sensing devices to govern use of data from the pools or stored on a device that governs use of storage capabilities of the device for a distributed ledger. Embodiments include training a model to determine what policies should be deployed in an industrial data collection system. Embodiments include a system for data collection in an industrial environment with a policy engine for deploying policy within the system and, optionally, self-organizing network coding for data transport. In example embodiments, a policy applies to how data will be presented in a multi-sensory interface, a heat map visual interface, or in an interface that operates with self-organized tuning of the interface layer.

[0386] As noted above, methods and systems are disclosed herein for on-device sensor fusion and data storage for industrial IoT devices, such as an industrial data collector, including self-organizing, remotely organized, or network-sensitive industrial data collectors, where data from multiple sensors is multiplexed at the device for storage of a fused data stream. Embodiments include a self-organizing marketplace that presents fused sensor data that is extracted from on-device storage of IoT devices. Embodiments include streaming fused sensor information from multiple industrial sensors and from an on-device data storage facility to a data pool. Embodiments include training a model to determine what data should be stored on a device in a data collection environment. Embodiments include a self-organizing swarm of industrial data collectors that organize among themselves to optimize data collection, where at least some of the data collectors have on-device storage of fused data from multiple sensors. Embodiments include storing distributed ledger information with fused sensor information on an industrial IoT device. Embodiments include a system for data collection with on-device sensor fusion, such as of industrial sensor data and, optionally, self-organizing network coding for data transport, where data structures are stored to support alternative, multi-sensory modes of presentation, visual heat map modes of presentation, and/or an interface that operates with self-organized tuning of the interface layer.

[0387] As noted above, methods and systems are disclosed herein for a self-organizing data marketplace for industrial IoT data, where available data elements are organized in the marketplace for consumption by consumers based on training a self-organizing facility with a training set and feedback from measures of marketplace success. Embodiments include organizing a set of data

pools in a self-organizing data marketplace based on utilization metrics for the data pools. Embodiments include training a model to determine pricing for data in a data marketplace. The data marketplace is fed with data streams from a self-organizing swarm of industrial data collectors, a set of industrial data collectors that have self-organizing storage, or self-organizing, network-sensitive, or remotely organized industrial data collectors. Embodiments include using a distributed ledger to store transactional data for a self-organizing marketplace for industrial IoT data. Embodiments include using self-organizing network coding for data transport to a marketplace for sensor data collected in industrial environments. Embodiments include providing a library of data structures suitable for presenting data in alternative, multi-sensory interface modes in a data marketplace, in heat map visualization, and/or in interfaces that operate with self-organized tuning of the interface layer.

[0388] As noted above, methods and systems are disclosed herein for self-organizing data pools such as those that self-organize based on utilization and/or yield metrics that may be tracked for a plurality of data pools. In example embodiments, the pools contain data from self-organizing data collectors. Embodiments include training a model to present the most valuable data in a data marketplace, where training is based on industry-specific measures of success. Embodiments include populating a set of self-organizing data pools with data from a self-organizing swarm of data collectors. Embodiments include using a distributed ledger to store transactional information for data that is deployed in data pools, where the distributed ledger is distributed across the data pools. Embodiments include populating a set of self-organizing data pools with data from a set of network-sensitive or remotely organized data collectors or a set of data collectors having self-organizing storage. Embodiments include a system for data collection in an industrial environment with self-organizing pools for data storage and self-organizing network coding for data transport, such as a system that includes a source data structure for supporting data presentation in a multi-sensory interface, in a heat map interface, and/or in an interface that operates with self-organized tuning of the interface layer.

[0389] As noted above, methods and systems are disclosed herein for training AI models based on industry-specific feedback, such as that reflects a measure of utilization, yield, or impact, where the AI model operates on sensor data from an industrial environment. Embodiments include training a swarm of data collectors, or data collectors, such as remotely organized, self-organizing, or network-sensitive data collectors, based on industry-specific feedback or network and industrial conditions in an industrial environment, such as to configure storage. Embodiments include training an AI model to identify and use available storage locations in an industrial environment for storing distributed ledger information. Embodiments include training a remote organizer for a remotely organized data collector based on industry-specific feedback measures. Embodiments include a system for data collection in an industrial environment with cloud-based training of a network coding model for organizing network coding for data transport or a facility that manages presentation of data in a multi-sensory interface, in a heat map interface, and/or in an interface that operates with self-organized tuning of the interface layer.

[0390] As noted above, methods and systems are disclosed herein for a self-organized swarm of industrial data collectors that organize among themselves to optimize data collection based on the capabilities and conditions of the members of the swarm. Embodiments include deploying distributed ledger data structures across a swarm of data. Data collectors may be network-sensitive data collectors configured for remote organization or have self-organizing storage. Systems for data collection in an industrial environment with a swarm can include a self-organizing network coding for data transport. Systems include swarms that relay information for use in a multi-sensory interface, in a heat map interface, and/or in an interface that operates with self-organized tuning of the interface layer.

[0391] As noted above, methods and systems are disclosed herein for an industrial IoT distributed ledger, including a distributed ledger supporting the tracking of transactions executed in an automated data marketplace for industrial IoT data. Embodiments include a self-organizing data collector that is configured to distribute collected information to a distributed ledger. Embodiments include a network-sensitive data collector that is configured to distribute collected information to a distributed ledger based on network conditions. Embodiments include a remotely organized data collector that is configured to distribute collected information to a distributed ledger based on intelligent, remote management of the distribution. Embodiments include a data collector with self-organizing local storage that is configured to distribute collected information to a distributed ledger. Embodiments include a system for data collection in an industrial environment using a distributed ledger for data storage and self-organizing network coding for data transport, wherein data storage is of a data structure supporting a haptic interface for data presentation, a heat map interface for data presentation, and/or an interface that operates with self-organized tuning of the interface layer.

[0392] As noted above, methods and systems are disclosed herein for a self-organizing collector, including a self-organizing, multi-sensor data collector that can optimize data collection, power and/or yield based on conditions in its environment, and is optionally responsive to remote organization. Embodiments include a self-organizing data collector that organizes at least in part based on network conditions. Embodiments include a self-organizing data collector with self-organizing storage for data collected in an industrial data collection environment. Embodiments include a system for data collection in an industrial environment with self-organizing data collection and self-organizing network coding for data transport. Embodiments include a system for data collection in an industrial environment with a self-organizing data collector that feeds a data structure supporting a haptic or multi-sensory wearable interface for data presentation, a heat map interface for data presentation, and/or an interface that operates with self-organized tuning of the interface layer.

[0393] As noted above, methods and systems are disclosed herein for a network-sensitive collector, including a network condition-sensitive, self-organizing, multi-sensor data collector that can optimize based on bandwidth, quality of service, pricing, and/or other network conditions. Embodiments include a remotely organized, network condition-sensitive universal data collector that can power up and down sensor interfaces based on need and/or conditions identified in an

industrial data collection environment, including network conditions. Embodiments include a network-condition sensitive data collector with self-organizing storage for data collected in an industrial data collection environment. Embodiments include a network-condition sensitive data collector with self-organizing network coding for data transport in an industrial data collection environment. Embodiments include a system for data collection in an industrial environment with a network-sensitive data collector that relays a data structure supporting a haptic wearable interface for data presentation, a heat map interface for data presentation, and/or an interface that operates with self-organized tuning of the interface layer.

[0394] As noted above, methods and systems are disclosed herein for a remotely organized universal data collector that can power up and down sensor interfaces based on need and/or conditions identified in an industrial data collection environment. Embodiments include a remotely organized universal data collector with self-organizing storage for data collected in an industrial data collection environment. Embodiments include a system for data collection in an industrial environment with remote control of data collection and self-organizing network coding for data transport. Embodiments include a remotely organized data collector for storing sensor data and delivering instructions for use of the data in a haptic or multi-sensory wearable interface, in a heat map visual interface, and/or in an interface that operates with self-organized tuning of the interface layer.

[0395] As noted above, methods and systems are disclosed herein for self-organizing storage for a multi-sensor data collector, including self-organizing storage for a multi-sensor data collector for industrial sensor data. Embodiments include a system for data collection in an industrial environment with self-organizing data storage and self-organizing network coding for data transport. Embodiments include a data collector with self-organizing storage for storing sensor data and instructions for translating the data for use in a haptic wearable interface, in a heat map presentation interface, and/or in an interface that operates with self-organized tuning of the interface layer.

[0396] As noted above, methods and systems are disclosed herein for self-organizing network coding for a multi-sensor data network, including self-organizing network coding for a data network that transports data from multiple sensors in an industrial data collection environment. The system includes a data structure supporting a haptic wearable interface for data presentation, a heat map interface for data presentation, and/or self-organized tuning of an interface layer for data presentation.

[0397] As noted above, methods and systems are disclosed herein for a haptic or multi-sensory user interface, including a wearable haptic or multi-sensory user interface for an industrial sensor data collector, with vibration, heat, electrical, and/or sound outputs. Embodiments include a wearable haptic user interface for conveying industrial state information from a data collector, with vibration, heat, electrical, and/or sound outputs. The wearable also has a visual presentation layer for presenting a heat map that indicates a parameter of the data. Embodiments include condition-sensitive, self-organized tuning of AR/VR interfaces and multi-sensory interfaces based on feedback metrics and/or training in industrial environments.

[0398] As noted above, methods and systems are disclosed herein for a presentation layer for AR/VR industrial glasses, where heat map elements are presented based on patterns and/or parameters in collected data. Embodiments include condition-sensitive, self-organized tuning of a heat map AR/VR interface based on feedback metrics and/or training in industrial environments. As noted above, methods and systems are disclosed herein for condition-sensitive, self-organized tuning of AR/VR interfaces based on feedback metrics and/or training in industrial environments.

[0399] The following illustrative clauses describe certain embodiments of the present disclosure. The data collection system mentioned in the following disclosure may be a local data collection system 102, a host processing system 112 (e.g., using a cloud platform), or a combination of a local system and a host system. In example embodiments, a data collection system or data collection and processing system is provided having the use of an analog crosspoint switch for collecting data having variable groups of analog sensor inputs and, in example embodiments, having IP front-end-end signal conditioning on a multiplexer for improved signal-to-noise ratio, multiplexer continuous monitoring alarming features, the use of distributed CPLD chips with a dedicated bus for logic control of multiple MUX and data acquisition sections, high-amperage input capability using solid state relays and design topology, power-down capability of at least one of an analog sensor channel and of a component board, unique electrostatic protection for trigger and vibration inputs, and/or precise voltage reference for A/D zero reference.

[0400] In example embodiments, a platform is provided having one or more of cloud-based, machine pattern recognition based on fusion of remote, analog industrial sensors, cloud-based, machine pattern analysis of state information from multiple analog industrial sensors to provide anticipated state information for an industrial system, a cloud-based policy automation engine for IoT, with creation, deployment, and management of IoT devices, on-device sensor fusion and data storage for industrial IoT devices, a self-organizing data marketplace for industrial IoT data, self-organization of data pools based on utilization and/or yield metrics, training AI models based on industry-specific feedback, a self-organized swarm of industrial data collectors, an IoT distributed ledger, a self-organizing collector, a network-sensitive collector, a remotely organized collector, a self-organizing storage for a multi-sensor data collector, a self-organizing network coding for multi-sensor data network, a wearable haptic user interface for an industrial sensor data collector, with vibration, heat, electrical, and/or sound outputs, heat maps displaying collected data for AR/VR, or automatically tuned AR/VR visualization of data collected by a data collector.

[0401] With regard to Figure 14, a range of existing data sensing and processing systems with industrial sensing, processing, and storage systems 4500 include a streaming data collector 4510 that may be configured to accept data in a range of formats as described herein. In example embodiments, the range of formats can include a data format A 4520, a data format B 4522, a data format C 4524, and a data format D 4528 that may be sourced from a range of sensors. Moreover, the range of sensors can include an instrument A 4540, an instrument B 4542, an instrument C 4544, and an instrument D 4548. The streaming data collector 4510 may be configured with processing capabilities that enable access to the individual formats while leveraging the streaming, routing, self-organizing storage, and other capabilities described herein.

[0402] Figure 15 depicts methods and systems 4600 for industrial machine sensor data streaming collection, processing, and storage that facilitate use of a streaming data collector 4610 to collect and obtain data from legacy instruments 4620 and streaming instruments 4622. Legacy instruments 4620 and their data methodologies may capture and provide data that is limited in scope, due to the legacy systems and acquisition procedures, such as existing data methodologies described herein, to a particular range of frequencies and the like. The streaming data collector 4610 may be configured to capture streaming instrument data 4632 as well as legacy instrument data 4630. The streaming data collector 4610 may also be configured to capture current streaming instruments 4620 and legacy instruments 4622 and sensors using current and legacy data methodologies. These embodiments may be useful in transition applications from the legacy instruments and processing to the streaming instruments and processing that may be current or desired instruments or methodologies. In example embodiments, the streaming data collector 4610 may be configured to process the legacy instrument data 4630 so that it can be stored compatibly with the streamed instrument data 4632. The streaming data collector 4610 may process or parse the streamed instrument data 4632 based on the legacy instrument data 4630 to produce at least one extraction of the streamed data 4642 that is compatible with the legacy instrument data 4630 that can be processed into translated legacy data 4640. In example embodiments, extracted data 4650 that can include extracted portions of translated legacy data 4652 and streamed data 4654 may be stored in a format that facilitates access and processing by legacy instrument data processing and further processing that can emulate legacy instrument data processing methods, and the like. In example embodiments, the portions of the translated legacy data 4652 may also be stored in a format that facilitates processing with different methods that can take advantage of the greater frequencies, resolution, and volume of data possible with a streaming instrument.

[0403] Figure 16 depicts alternate embodiments descriptive of methods and systems 4700 for industrial machine sensor data streaming, collection, processing, and storage that facilitate integration of legacy instruments and processing. In example embodiments, a streaming data collector 4710 may be connected with an industrial machine 4712 and may include a plurality of sensors, such as streaming sensors 4720 and 4722 that may be configured to sense aspects of the industrial machine 4712 associated with at least one moving part of the machine 4712. The sensors 4720 and 4722 (or more) may communicate with one or more streaming devices 4740 that may facilitate streaming data from one or more of the sensors to the streaming data collector 4710. In example embodiments, the industrial machine 4712 may also interface with or include one or more legacy instruments 4730 that may capture data associated with one or more moving parts of the industrial machine 4712 and store that data into a legacy data storage facility 4732.

[0404] In example embodiments, a frequency and/or resolution detection facility 4742 may be configured to facilitate detecting information about legacy instrument sourced data, such as a frequency range of the data or a resolution of the data, and the like. The detection facility 4742 may operate on data directly from the legacy instruments 4730 or from data stored in a legacy storage facility 4732. The detection facility 4742 may communicate information detected about the legacy instruments 4730, its sourced data, and its stored data 4732, or the like to the streaming data

collector 4710. Alternatively, the detection facility 4742 may access information, such as information about frequency ranges, resolution, and the like that characterizes the sourced data from the legacy instrument 4730 and/or may be accessed from a portion of the legacy storage facility 4732.

[0405] In example embodiments, the streaming data collector 4710 may be configured with one or more automatic processors, algorithms, and/or other data methodologies to match up information captured by the one or more legacy instruments 4730 with a portion of data being provided by the one or more streaming devices 4740 from the one or more industrial machines 4712. Data from streaming devices 4740 may include a wider range of frequencies and resolutions than the sourced data of legacy instruments 4730 and, therefore, filtering and other such functions can be implemented to extract data from the streaming devices 4740 that corresponds to the sourced data of the legacy instruments 4730 in aspects such as frequency range, resolution, and the like. In example embodiments, the configured streaming data collector 4710 may produce a plurality of streams of data, including a stream of data that may correspond to the stream of data from the streaming device 4740 and a separate stream of data that is compatible, in some aspects, with the legacy instrument sourced data and the infrastructure to ingest and automatically process it. Alternatively, the streaming data collector 4710 may output data in modes other than as a stream, such as batches, aggregations, summaries, and the like.

[0406] Configured streaming data collector 4710 may communicate with a stream storage facility 4764 for storing at least one of the data outputs from the streaming device 4710 and data extracted therefrom that may be compatible, in some aspects, with the sourced data of the legacy instruments 4730. A legacy compatible output of the configured streaming data collector 4710 may also be provided to a format adaptor facility 4748, 4760 that may configure, adapt, reformat, and make other adjustments to the legacy compatible data so that it can be stored in a legacy compatible storage facility 4762 so that legacy processing facilities 4744 may execute data processing methods on data in the legacy compatible storage facility 4762 and the like that are configured to process the sourced data of the legacy instruments 4730. In embodiments in which legacy compatible data is stored in the stream storage facility 4764, legacy processing facility 4744 may also automatically process this data after optionally being processed by format adaptor 4760. By arranging the data collection, streaming, processing, formatting, and storage elements to provide data in a format that is fully compatible with legacy instrument sourced data, transition from a legacy system can be simplified, and the sourced data from legacy instruments can be easily compared to newly acquired data (with more content) without losing the legacy value of the sourced data from the legacy instruments 4730.

[0407] Figure 17 depicts alternate embodiments of the methods and systems 4800 described herein for industrial machine sensor data streaming, collection, processing, and storage that may be compatible with legacy instrument data collection and processing. In example embodiments, processing industrial machine sensed data may be accomplished in a variety of ways including aligning legacy and streaming sources of data, such as by aligning stored legacy and streaming data; aligning stored legacy data with a stream of sensed data; and aligning legacy and streamed

data as it is being collected. In example embodiments, an industrial machine 4810 may include, communicate with, or be integrated with one or more stream data sensors 4820 that may sense aspects of the industrial machine 4810 such as aspects of one or more moving parts of the machine. The industrial machine 4810 may also communicate with, include, or be integrated with one or more legacy data sensors 4830 that may sense similar aspects of the industrial machine 4810. In example embodiments, the one or more legacy data sensors 4830 may provide sensed data to one or more legacy data collectors 4840. The stream data sensors 4820 may produce an output that encompasses all aspects of (i.e., a richer signal) and is compatible with sensed data from the legacy data sensors 4830. The stream data sensors 4820 may provide compatible data to the legacy data collector 4840. By mimicking the legacy data sensors 4830 or their data streams, the stream data sensors 4820 may replace (or serve as suitable duplicate for) one or more legacy data sensors, such as during an upgrade of the sensing and processing system of an industrial machine. Frequency range, resolution, and the like may be mimicked by the stream data so as to ensure that all forms of legacy data are captured or can be derived from the stream data. In example embodiments, format conversion, if needed, can also be performed by the stream data sensors 4820. The stream data sensors 4820 may also produce an alternate data stream that is suitable for collection by the stream data collector 4850. In example embodiments, such an alternate data stream may be a superset of the legacy data sensor data in at least one or more of: frequency range, resolution, duration of sensing the data, and the like.

[0408] In example embodiments, an industrial machine sensed data processing facility 4860 may execute a wide range of sensed data processing methods, some of which may be compatible with the data from legacy data sensors 4830 and may produce outputs that may meet legacy sensed data processing requirements. To facilitate use of a wide range of data processing capabilities of processing facility 4860, legacy and stream data may need to be aligned so that a compatible portion of stream data may be extracted for processing with legacy compatible methods and the like. In example embodiments, Figure 17 depicts three different techniques for aligning stream data to legacy data. A first alignment methodology 4862 includes aligning legacy data output by the legacy data collector 4840 with stream data output by the stream data collector 4850. As data is provided by the legacy data collector 4840, aspects of the data may be detected, such as resolution, frequency, duration, and the like, and may be used as control for a processing method that identifies portions of a stream of data from the stream data collector 4850 that are purposely compatible with the legacy data. The processing facility 4860 may apply one or more legacy compatible methods on the identified portions of the stream data to extract data that can be easily compared to or referenced against the legacy data.

[0409] In example embodiments, a second alignment methodology 4864 may involve aligning streaming data with data from a legacy storage facility 4882. In example embodiments, a third alignment methodology 4868 may involve aligning stored stream data from a stream storage facility 4884 with legacy data from the legacy data storage facility 4882. In each of the methodologies 4862, 4864, 4868, alignment data may be determined by processing the legacy data to detect aspects such as resolution, duration, frequency range, and the like. Alternatively,

alignment may be performed by an alignment facility, such as facilities using methodologies 4862, 4864, 4868 that may receive or may be configured with legacy data descriptive information such as legacy frequency range, duration, resolution, and the like.

[0410] In example embodiments, an industrial machine sensing data processing facility 4860 may have access to legacy compatible methods and algorithms that may be stored in a legacy data methodology storage facility 4880. These methodologies, algorithms, or other data in the legacy algorithm storage facility 4880 may also be a source of alignment information that could be communicated by the industrial machine sensed data processing facility 4860 to the various alignment facilities having methodologies 4862, 4864, 4868. By having access to legacy compatible algorithms and methodologies, the data processing facility 4860 may facilitate processing legacy data, streamed data that is compatible with legacy data, or portions of streamed data that represent the legacy data to produce legacy compatible analytics.

[0411] In example embodiments, the data processing facility 4860 may execute a wide range of other sensed data processing methods, such as wavelet derivations and the like, to produce streamed data analytics 4892. In example embodiments, the streaming data collector 102, 4510, 4610, 4710 (Figures 3, 6, 14, 15, 16) or data processing facility 4860 may include portable algorithms, methodologies, and inputs that may be defined and extracted from data streams. In many examples, a user or enterprise may already have existing and effective methods related to analyzing specific pieces of machinery and assets. These existing methods could be imported into the configured streaming data collector 102, 4510, 4610, 4710 or the data processing facility 4860 as portable algorithms or methodologies. Data processing, such as described herein for the configured streaming data collector 102, 4510, 4610, 4710 may also match an algorithm or methodology to a situation, then extract data from a stream to match to the data methodology from the legacy acquisition or legacy acquisition techniques. In example embodiments, the streaming data collector 102, 4510, 4610, 4710 may be compatible with many types of systems and may be compatible with systems having varying degrees of criticality.

[0412] Exemplary industrial machine deployments of the methods and systems described herein are now described. An industrial machine may be a gas compressor. In an example, a gas compressor may operate an oil pump on a very large turbo machine, such as a very large turbo machine that includes 10,000 HP motors. The oil pump may be a highly critical system as its failure could cause an entire plant to shut down. The gas compressor in this example may run four stages at a very high frequency, such as 36,000 RPM, and may include tilt pad bearings that ride on an oil film. The oil pump in this example may have roller bearings, such that if an anticipated failure is not being picked up by a user, the oil pump may stop running, and the entire turbo machine would fail. Continuing with this example, the streaming data collector 102, 4510, 4610, 4710 may collect data related to vibrations, such as casing vibration and proximity probe vibration. Other bearings industrial machine examples may include generators, power plants, boiler feed pumps, fans, forced draft fans, induced draft fans, and the like. The streaming data collector 102, 4510, 4610, 4710 for a bearings system used in the industrial gas industry may support predictive analysis

on the motors, such as that performed by model-based expert systems – for example, using voltage, current, and vibration as analysis metrics.

[0413] Another exemplary industrial machine deployment may be a motor and the streaming data collector 102, 4510, 4610, 4710 that may assist in the analysis of a motor by collecting voltage and current data on the motor, for example.

[0414] Yet another exemplary industrial machine deployment may include oil quality sensing. An industrial machine may conduct oil analysis, and the streaming data collector 102, 4510, 4610, 4710 may assist in searching for fragments of metal in oil, for example.

[0415] The methods and systems described herein may also be used in combination with model-based systems. Model-based systems may integrate with proximity probes. Proximity probes may be used to sense problems with machinery and shut machinery down due to sensed problems. A model-based system integrated with proximity probes may measure a peak waveform and send a signal that shuts down machinery based on the peak waveform measurement.

[0416] Enterprises that operate industrial machines may operate in many diverse industries. These industries may include industries that operate manufacturing lines, provide computing infrastructure, support financial services, provide HVAC equipment, and the like. These industries may be highly sensitive to lost operating time and the cost incurred due to lost operating time. HVAC equipment enterprises in particular may be concerned with data related to ultrasound, vibration, IR, and the like, and may get much more information about machine performance related to these metrics using the methods and systems of industrial machine sensed data streaming collection than from legacy systems.

[0417] Methods and systems described herein for industrial machine sensor data streaming, collection, processing, and storage may be configured to operate and integrate with existing data collection, processing and storage systems and may include a method for capturing a plurality of streams of sensed data from sensors deployed to monitor aspects of an industrial machine associated with at least one moving part of the machine; at least one of the streams containing a plurality of frequencies of data. The method may include identifying a subset of data in at least one of the multiple streams that corresponds to data representing at least one predefined frequency. The at least one predefined frequency is represented by a set of data collected from alternate sensors deployed to monitor aspects of the industrial machine associated with the at least one moving part of the machine. The method may further include processing the identified data with a data processing facility that processes the identified data with data methodologies configured to be applied to the set of data collected from alternate sensors. Lastly, the method may include storing the at least one of the streams of data, the identified subset of data, and a result of processing the identified data in an electronic data set.

[0418] The methods and systems may include a method for applying data captured from sensors deployed to monitor aspects of an industrial machine associated with at least one moving part of the machine, the data captured with predefined lines of resolution covering a predefined frequency range, to a frequency matching facility that identifies a subset of data streamed from other sensors deployed to monitor aspects of the industrial machine associated with at least one moving part of

the machine, the streamed data comprising a plurality of lines of resolution and frequency ranges, the subset of data identified corresponding to the lines of resolution and predefined frequency range. This method may include storing the subset of data in an electronic data record in a format that corresponds to a format of the data captured with predefined lines of resolution, and signaling to a data processing facility the presence of the stored subset of data. This method may optionally include processing the subset of data with at least one of algorithms, methodologies, models, and pattern recognizers that corresponds to algorithms, methodologies, models, and pattern recognizers associated with processing the data captured with predefined lines of resolution covering a predefined frequency range.

[0419] The methods and systems may include a method for identifying a subset of streamed sensor data. The sensor data is captured from sensors deployed to monitor aspects of an industrial machine associated with at least one moving part of the machine. The subset of streamed sensor data is at predefined lines of resolution for a predefined frequency range. The method includes establishing a first logical route for communicating electronically between a first computing facility performing the identifying and a second computing facility. The identified subset of the streamed sensor data is communicated exclusively over the established first logical route when communicating the subset of streamed sensor data from the first facility to the second facility. This method may further include establishing a second logical route for communicating electronically between the first computing facility and the second computing facility for at least one portion of the streamed sensor data that is not the identified subset. This method may further include establishing a third logical route for communicating electronically between the first computing facility and the second computing facility for at least one portion of the streamed sensor data that includes the identified subset and at least one other portion of the data not represented by the identified subset.

[0420] The methods and systems may include a first data sensing and processing system that captures first data from a first set of sensors deployed to monitor aspects of an industrial machine associated with at least one moving part of the machine, the first data covering a set of lines of resolution and a frequency range. This system may include a second data sensing and processing system that captures and streams a second set of data from a second set of sensors deployed to monitor aspects of the industrial machine associated with at least one moving part of the machine, the second data covering a plurality of lines of resolution that includes the set of lines of resolution and a plurality of frequencies that includes the frequency range. The system may enable: (1) selecting a portion of the second data that corresponds to the set of lines of resolution and the frequency range of the first data; and (2) processing the selected portion of the second data with the first data sensing and processing system.

[0421] The methods and systems may include a method for automatically processing a portion of a stream of sensed data. The sensed data received from a first set of sensors deployed to monitor aspects of an industrial machine associated with at least one moving part of the machine in response to an electronic data structure that facilitates extracting a subset of the stream of sensed data that corresponds to a set of sensed data received from a second set of sensors deployed to monitor the aspects of the industrial machine associated with the at least one moving part of the machine. The

set of sensed data is constrained to a frequency range. The stream of sensed data includes a range of frequencies that exceeds the frequency range of the set of sensed data. The processing comprises executing data methodologies on a portion of the stream of sensed data that is constrained to the frequency range of the set of sensed data. The data methodologies are configured to process the set of sensed data.

[0422] The methods and systems may include a method for receiving first data from sensors deployed to monitor aspects of an industrial machine associated with at least one moving part of the machine. This method may further include: (1) detecting at least one of a frequency range and lines of resolution represented by the first data, and (2) receiving a stream of data from sensors deployed to monitor the aspects of the industrial machine associated with the at least one moving part of the machine. The stream of data includes: a plurality of frequency ranges and a plurality of lines of resolution that exceeds the frequency range and the lines of resolution represented by the first data; extracting a set of data from the stream of data that corresponds to at least one of the frequency range and the lines of resolution represented by the first data; and processing the extracted set of data with a data processing method that is configured to process data within the frequency range and within the lines of resolution of the first data.

[0423] The methods and systems disclosed herein may include, connect to, or be integrated with a data acquisition instrument and in the many embodiments, Figure 18 shows methods and systems 5000 that includes a data acquisition (DAQ) streaming instrument 5002 also known as an SDAQ. In example embodiments, output from sensors 5010, 5012, 5014 may be of various types including vibration, temperature, pressure, ultrasound and so on. In my many examples, one of the sensors may be used. In further examples, many of the sensors may be used and their signals may be used individually or in predetermined combinations and/or at predetermined intervals, circumstances, setups, and the like.

[0424] In example embodiments, the output signals from the sensors 5010, 5012, 5014 may be fed into instrument inputs 5020, 5022, 5024 of the DAQ instrument 5002 and may be configured with additional streaming capabilities 5028. By way of these many examples, the output signals from the sensors 5010, 5012, 5014, or more as applicable, may be conditioned as an analog signal before digitization with respect to at least scaling and filtering. The signals may then be digitized by an analog-to-digital converter 5030. The signals received from all relevant channels (i.e., one or more channels are switched on manually, by alarm, by route, and the like) may be simultaneously sampled at a predetermined rate sufficient to perform the maximum desired frequency analysis that may be adjusted and readjusted as needed or otherwise held constant to ensure compatibility or conformance with other relevant datasets. In example embodiments, the signals are sampled for a relatively long time and gap-free as one continuous stream so as to enable further post-processing at lower sampling rates with sufficient individual sampling.

[0425] In example embodiments, data may be streamed from a collection of points and then the next set of data may be collected from additional points according to a prescribed sequence, route, path, or the like. In many examples, the sensors 5010, 5012, 5014 or more may be moved to the next location according to the prescribed sequence, route, pre-arranged configurations, or the like.

In certain examples, not all of the sensors 5010, 5012, 5014 may move and therefore some may remain fixed in place and used for detection of reference phase or the like.

[0426] In example embodiments, a multiplex (mux) 5032 may be used to switch to the next collection of points, to a mixture of the two methods or collection patterns that may be combined, other predetermined routes, and the like. The multiplexer 5032 may be stackable so as to be laddered and effectively accept more channels than the DAQ instrument 5002 provides. In examples, the DAQ instrument 5002 may provide eight channels while the multiplexer 5032 may be stacked to supply 32 channels. Further variations are possible with one more multiplexers. In example embodiments, the multiplexer 5032 may be fed into the DAQ instrument 5002 through an instrument input 5034. In example embodiments, the DAQ instrument 5002 may include a controller 5038 that may take the form of an onboard controller, a PC, other connected devices, network based services, and combinations thereof.

[0427] In example embodiments, the sequence and panel conditions used to govern the data collection process may be obtained from the multimedia probe (MMP) and probe control, sequence and analytical (PCSA) information store 5040. In example embodiments, the information store 5040 may be onboard the DAQ instrument 5002. In example embodiments, contents of the information store 5040 may be obtained through a cloud network facility, from other DAQ instruments, from other connected devices, from the machine being sensed, other relevant sources, and combinations thereof. In example embodiments, the information store 5040 may include such items as the hierarchical structural relationships of the machine, e.g., a machine contains predetermined pieces of equipment, each of which may contain one or more shafts and each of those shafts may have multiple associated bearings. Each of those types of bearings may be monitored by specific types of transducers or probes, according to one or more specific prescribed sequences (paths, routes, and the like) and with one or more specific panel conditions that may be set on the one or more DAQ instruments 5002. By way of this example, the panel conditions may include hardware specific switch settings or other collection parameters. In many examples, collection parameters include but are not limited to a sampling rate, AC/DC coupling, voltage range and gain, integration, high and low pass filtering, anti-aliasing filtering, ICP™ transducers and other integrated-circuit piezoelectric transducers, 4-20 mA loop sensors, and the like. In example embodiments, the information store 5040 may also include machinery specific features that may be important for proper analysis such as gear teeth for a gear, number blades in a pump impeller, number of motor rotor bars, bearing specific parameters necessary for calculating bearing frequencies, revolution per minutes information of all rotating elements and multiples of those RPM ranges, and the like. Information in the information store may also be used to extract stream data 5050 for permanent storage.

[0428] Based on directions from the DAQ API software 5052, digitized waveforms may be uploaded using DAQ driver services 5054 of a driver onboard the DAQ instrument 5002. In example embodiments, data may then be fed into a raw data server 5058 which may store the stream data 5050 in a stream data repository 5060. In example embodiments, this data storage area is typically meant for storage until the data is copied off of the DAQ instrument 5002 and verified.

The DAQ API 5052 may also direct the local data control application 5062 to extract and process the recently obtained stream data 5050 and convert it to the same or lower sampling rates of sufficient length to affect one or more desired resolutions. By way of these examples, this data may be converted to spectra, averaged, and processed in a variety of ways and stored, at least temporarily, as extracted/processed (EP) data 5064. It will be appreciated in light of the disclosure that legacy data may require its own sampling rates and resolution to ensure compatibility and often this sampling rate may not be integer proportional to the acquired sampling rate. It will also be appreciated in light of the disclosure that this may be especially relevant for order-sampled data whose sampling frequency is related directly to an external frequency (typically the running speed of the machine or its local componentry) rather than the more-standard sampling rates employed by the internal crystals, clock functions, or the like of the DAQ instrument (e.g., values of Fmax of 100, 200, 500, 1K, 2K, 5K, 10K, 20K, and so on).

[0429] In example embodiments, the extract/process (EP) align module 5068 of the local data control application 5062 may be able to fractionally adjust the sampling rates to these non-integer ratio rates satisfying an important requirement for making data compatible with legacy systems. In example embodiments, fractional rates may also be converted to integer ratio rates more readily because the length of the data to be processed may be adjustable. It will be appreciated in light of the disclosure that if the data was not streamed and just stored as spectra with the standard or predetermined Fmax, it may be impossible in certain situations to convert it retroactively and accurately to the order-sampled data. It will also be appreciated in light of the disclosure that internal identification issues may also need to be reconciled. In many examples, stream data may be converted to the proper sampling rate and resolution as described and stored (albeit temporarily) in an EP legacy data repository 5070 to ensure compatibility with legacy data.

[0430] To support legacy data identification issues, a user input module 5072 is shown in many embodiments should there be no automated process (whether partially or wholly) for identification translation. In such examples, one or more legacy systems (i.e., pre-existing data acquisition) may be characterized in that the data to be imported is in a fully standardized format such as a Mimosa™ format, and other similar formats. Moreover, sufficient indentation of the legacy data and/or the one or more machines from which the legacy data was produced may be required in the completion of an identification mapping table 5074 to associate and link a portion of the legacy data to a portion of the newly acquired streamed data 5050. In many examples, the end user and/or legacy vendor may be able to supply sufficient information to complete at least a portion of a functioning identification (ID) mapping table 5074 and therefore may provide the necessary database schema for the raw data of the legacy system to be used for comparison, analysis, and manipulation of newly streamed data 5050.

[0431] In example embodiments, the local data control application 5062 may also direct streaming data as well as extracted/processed (EP) data to a cloud network facility 5080 via wired or wireless transmission. From the cloud network facility 5080 other devices may access, receive, and maintain data including the data from a master raw data server (MRDS) 5082. The movement,

distribution, storage, and retrieval of data remote to the DAQ instrument 5002 may be coordinated by the cloud data management services ("CDMS") 5084.

[0432] Figure 19 shows additional methods and systems that include the DAQ instrument 5002 accessing related cloud based services. In example embodiments, the DAQ API 5052 may control the data collection process as well as its sequence. By way of these examples, the DAQ API 5052 may provide the capability for editing processes, viewing plots of the data, controlling the processing of that data, viewing the output data in all its myriad forms, analyzing this data including expert analysis, and communicating with external devices via the local data control application 5062 and with the CDMS 5084 via the cloud network facility 5080. In example embodiments, the DAQ API 5052 may also govern the movement of data, its filtering, as well as many other housekeeping functions.

[0433] In example embodiments, an expert analysis module 5100 may generate reports 5102 that may use machine or measurement point specific information from the information store 5040 to analyze the stream data 5050 using a stream data analyzer module 5104 and the local data control application 5062 with the extract/process ("EP") align module 5068. In example embodiments, the expert analysis module 5100 may generate new alarms or ingest alarm settings into an alarms module 5108 that is relevant to the stream data 5050. In example embodiments, the stream data analyzer module 5104 may provide a manual or automated mechanism for extracting meaningful information from the stream data 5050 in a variety of plotting and report formats. In example embodiments, a supervisory control of the expert analysis module 5100 is provided by the DAQ API 5052. In further examples, the expert analysis module 5100 may be supplied (wholly or partially) via the cloud network facility 5080. In many examples, the expert analysis module 5100 via the cloud may be used rather than a locally-deployed expert analysis module 5100 for various reasons such as using the most up-to-date software version, more processing capability, a bigger volume of historical data to reference, and so on. In many examples, it may be important that the expert analysis module 5100 be available when an internet connection cannot be established so having this redundancy may be crucial for seamless and time efficient operation. Toward that end, many of the modular software applications and databases available to the DAQ instrument 5002 where applicable may be implemented with system component redundancy to provide operational robustness to provide connectivity to cloud services when needed but also operate successfully in isolated scenarios where connectivity is not available and sometime not available purposefully to increase security and the like.

[0434] In example embodiments, the DAQ instrument acquisition may require a real time operating system ("RTOS") for the hardware especially for streamed gap-free data that is acquired by a PC. In some instances, the requirement for a RTOS may result in (or may require) expensive custom hardware and software capable of running such a system. In many embodiments, such expensive custom hardware and software may be avoided and an RTOS may be effectively and sufficiently implemented using a standard Windows™ operating systems or similar environments including the system interrupts in the procedural flow of a dedicated application included in such operating systems.

[0435] The methods and systems disclosed herein may include, connect to, or be integrated with one or more DAQ instruments and in the many embodiments, Figure 20 shows methods and systems 5150 that include the DAQ instrument 5002 (also known as a streaming DAQ or an SDAQ). In example embodiments, the DAQ instrument 5002 may effectively and sufficiently implement an RTOS using standard windows operating system (or other similar personal computing systems) that may include a software driver configured with a First In, First Out (FIFO) memory area 5152. The FIFO memory area 5152 may be maintained and hold information for a sufficient amount of time to handle a worst-case interrupt that it may face from the local operating system to effectively provide the RTOS. In many examples, configurations on a local personal computer or connected device may be maintained to minimize operating system interrupts. To support this, the configurations may be maintained, controlled, or adjusted to eliminate (or be isolated from) any exposure to extreme environments where operating system interrupts may become an issue. In example embodiments, the DAQ instrument 5002 may produce a notification, alarm, message, or the like to notify a user when any gap errors are detected. In these many examples, such errors may be shown to be rare and even if they occur, the data may be adjusted knowing when they occurred should such a situation arise.

[0436] In example embodiments, the DAQ instrument 5002 may maintain a sufficiently large FIFO memory area 5152 that may buffer the incoming data so as to be not affected by operating system interrupts when acquiring data. It will be appreciated in light of the disclosure that the predetermined size of the FIFO memory area 5152 may be based on operating system interrupts that may include Windows system and application functions such as the writing of data to Disk or SSD, plotting, GUI interactions and standard Windows tasks, low-level driver tasks such as servicing the DAQ hardware and retrieving the data in bursts, and the like.

[0437] In example embodiments, the computer, controller, connected device or the like that may be included in the DAQ instrument 5002 may be configured to acquire data from the one or more hardware devices over a USB port, firewire, ethernet, or the like. In example embodiments, the DAQ driver services 5054 may be configured to have data delivered to it periodically so as to facilitate providing a channel specific FIFO memory buffer that may be configured to not miss data, i.e., it is gap-free. In example embodiments, the DAQ driver services 5054 may be configured so as to maintain an even larger (than the device) channel specific FIFO area 5152 that it fills with new data obtained from the device. In example embodiments, the DAQ driver services 5054 may be configured to employ a further process in that the raw data server 5058 may take data from the FIFO 5110 and may write it as a contiguous stream to non-volatile storage areas such as the stream data repository 5060 that may be configured as one or more disk drives, SSDs, or the like. In example embodiments, the FIFO 5110 may be configured to include a starting and stopping marker or pointer to mark where the latest most current stream was written. By way of these examples, a FIFO end marker 5114 may be configured to mark the end of the most current data until it reaches the end of the spooler and then wraps around constantly cycling around. In these examples, there is always one megabyte (or other configured capacities) of the most current data available in the FIFO 5110 once the spooler fills up. It will be appreciated in light of the disclosure that further

configurations of the FIFO memory area may be employed. In example embodiments, the DAQ driver services 5054 may be configured to use the DAQ API 5052 to pipe the most recent data to a high-level application for processing, graphing and analysis purposes. In some examples, it is not required that this data be gap-free but even in these instances, it is helpful to identify and mark the gaps in the data. Moreover, these data updates may be configured to be frequent enough so that the user would perceive the data as live. In the many embodiments, the raw data is flushed to non-volatile storage without a gap at least for the prescribed amount of time and examples of the prescribed amount of time may be about thirty seconds to over four hours. It will be appreciated in light of the disclosure that many pieces of equipment and their components may contribute to the relative needed duration of the stream of gap-free data and those durations may be over four hours when relatively low speeds are present in large numbers, when non-periodic transient activity is occurring on a relatively long time frame, when duty cycle only permits operation in relevant ranges for restricted durations and the like.

[0438] With reference to Figure 19, the stream data analyzer module 5104 may provide for the manual or extraction of information from the data stream in a variety of plotting and report formats. In example embodiments, resampling, filtering (including anti-aliasing), transfer functions, spectrum analysis, enveloping, averaging, peak detection functionality, as well as a host of other signal processing tools, may be available for the analyst to analyze the stream data and to generate a very large array of snapshots. It will be appreciated in light of the disclosure that much larger arrays of snapshots are created than ever would have been possible by scheduling the collection of snapshots beforehand, i.e., during the initial data acquisition for the measurement point in question.

[0439] It will be appreciated in light of the disclosure that the sampling rates of vibration data of up to 100 kHz (or higher in some scenarios) may be utilized for non-vibration sensors as well. In doing so, it will further be appreciated in light of the disclosure that stream data in such durations at these sampling rates may uncover new patterns to be analyzed due in no small part that many of these types of sensors have not been utilized in this manner. It will also be appreciated in light of the disclosure that different sensors used in machinery condition monitoring may provide measurements more akin to static levels rather than fast-acting dynamic signals. In some cases, faster response time transducers may have to be used prior to achieving the faster sampling rates.

[0440] In many embodiments, sensors may have a relatively static output such as temperature, pressure, or flow but may still be analyzed with the dynamic signal processing system and methodologies as disclosed herein. It will be appreciated in light of the disclosure that the time scale, in many examples, may be slowed down. In many examples, a collection of temperature readings collected approximately every minute for over two weeks may be analyzed for their variation solely or in collaboration or in fusion with other relevant sensors. By way of these examples, the direct current level or average level may be omitted from all the readings (e.g., by subtraction) and the resulting delta measurements may be processed (e.g., through a Fourier transform). From these examples, resulting spectral lines may correlate to specific machinery behavior or other symptoms present in industrial system processes. In further examples, other techniques include enveloping that may look for modulation, wavelets that may look for spectral

patterns that last only for a short time (e.g., bursts), cross-channel analysis to look for correlations with other sensors including vibration, and the like.

[0441] Figure 21 shows a DAQ instrument 5400 that may be integrated with one or more analog sensors 5402 and endpoint nodes 5404 to provide a streaming sensor 5410 or smart sensors that may take in analog signals and then process and digitize them, and then transmit them to one or more external monitoring systems 5412 in the many embodiments that may be connected to, interfacing with, or integrated with the methods and systems disclosed herein. The monitoring system 5412 may include a streaming hub server 5420 that may communicate with the CDMS 5084. In example embodiments, the CDMS 5084 may contact, use, and integrate with cloud data 5430 and cloud services 5432 that may be accessible through one or more cloud network facilities 5080. In example embodiments, the streaming hub server 5420 may connect with another streaming sensor 5440 that may include a DAQ instrument 5442, an endpoint node 5444, and the one or more analog sensors such as analog sensor 5448. The steaming hub server 5420 may connect with other streaming sensors such as the streaming sensor 5460 that may include a DAQ instrument 5462, an endpoint node 5464, and the one or more analog sensors such as analog sensor 5468.

[0442] In example embodiments, there may be additional streaming hub servers such as the steaming hub server 5480 that may connect with other streaming sensors such as the streaming sensor 5490 that may include a DAQ instrument 5492, an endpoint node 5494, and the one or more analog sensors such as analog sensor 5498. In example embodiments, the streaming hub server 5480 may also connect with other streaming sensors such as the streaming sensor 5500 that may include a DAQ instrument 5502, an endpoint node 5504, and the one or more analog sensors such as analog sensor 5508. In example embodiments, the transmission may include averaged overall levels and in other examples may include dynamic signal sampled at a prescribed and/or fixed rate. In example embodiments, the streaming sensors 5410, 5440, 5460, 5490, and 5500 may be configured to acquire analog signals and then apply signal conditioning to those analog signals including coupling, averaging, integrating, differentiating, scaling, filtering of various kinds, and the like. The streaming sensors 5410, 5440, 5460, 5490, and 5500 may be configured to digitize the analog signals at an acceptable rate and resolution (number of bits) and to process further the digitized signal when required. The streaming sensors 5410, 5440, 5460, 5490, and 5500 may be configured to transmit the digitized signals at pre-determined, adjustable, and re-adjustable rates. In example embodiments, the streaming sensors 5410, 5440, 5460, 5490, and 5500 are configured to acquire, digitize, process, and transmit data at a sufficient effective rate so that a relatively consistent stream of data may be maintained for a suitable amount of time so that a large number of effective analyses may be shown to be possible. In the many embodiments, there would be no gaps in the data stream and the length of data should be relatively long, ideally for an unlimited amount of time, although practical considerations typically require ending the stream. It will be appreciated in light of the disclosure that this long duration data stream with effectively no gap in the stream is in contrast to the more commonly used burst collection where data is collected for a relatively short period of time (i.e., a short burst of collection), followed by a pause, and then perhaps another burst collection and so on. In the commonly used collections of data collected over

noncontiguous bursts, data would be collected at a slow rate for low frequency analysis and high frequency for high frequency analysis. In many embodiments of the present disclosure, in contrast, the streaming data is being collected (i) once, (ii) at the highest useful and possible sampling rate, and (iii) for a long enough time that low frequency analysis may be performed as well as high frequency. To facilitate the collection of the streaming data, enough storage memory must be available on the one or more streaming sensors such as the streaming sensors 5410, 5440, 5460, 5490, 5500 so that new data may be off-loaded externally to another system before the memory overflows. In example embodiments, data in this memory would be stored into and accessed from "First-In, First-Out" ("FIFO") mode. In these examples, the memory with a FIFO area may be a dual port so that the sensor controller may write to one part of it while the external system reads from a different part. In example embodiments, data flow traffic may be managed by semaphore logic.

[0443] It will be appreciated in light of the disclosure that vibration transducers that are larger in mass will have a lower linear frequency response range because the natural resonance of the probe is inversely related to the square root of the mass and will be lowered. Toward that end, a resonant response is inherently non-linear and so a transducer with a lower natural frequency will have a narrower linear passband frequency response. It will also be appreciated in light of the disclosure that above the natural frequency the amplitude response of the sensor will taper off to negligible levels rendering it even more unusable. With that in mind, high frequency accelerometers, for this reason, tend to be quite small in mass, to the order of half of a gram. It will also be appreciated in light of the disclosure that adding the required signal processing and digitizing electronics required for streaming may, in certain situations, render the sensors incapable in many instances of measuring high-frequency activity.

[0444] In example embodiments, streaming hubs such as the streaming hubs 5420, 5480 may effectively move the electronics required for streaming to an external hub via cable. It will be appreciated in light of the disclosure that the streaming hubs may be located virtually next to the streaming sensors or up to a distance supported by the electronic driving capability of the hub. In instances where an internet cache protocol ("ICP") is used, the distance supported by the electronic driving capability of the hub would be anywhere from 100 to 1000 feet (30.5 to 305 meters) based on desired frequency response, cable capacitance, and the like. In example embodiments, the streaming hubs may be positioned in a location convenient for receiving power as well as connecting to a network (be it LAN or WAN). In example embodiments, other power options would include solar, thermal as well as energy harvesting. Transfer between the streaming sensors and any external systems may be wireless or wired and may include such standard communication technologies as 802.11 and 900 MHz wireless systems, Ethernet, USB, firewire and so on.

[0445] With reference to Figure 18, the many examples of the DAQ instrument 5002 include embodiments where data that may be uploaded from the local data control application 5062 to the master raw data server ("MRDS") 5082. In example embodiments, information in the multimedia probe ("MMP") and probe control, sequence and analytical ("PCSA") information store 5040 may also be downloaded from the MRDS 5082 down to the DAQ instrument 5002. Further details of

the MRDS 5082 are shown in Figure 22 including embodiments where data may be transferred to the MRDS 5082 from the DAQ instrument 5002 via a wired or wireless network, or through connection to one or more portable media, drive, other network connections, or the like. In example embodiments, the DAQ instrument 5002 may be configured to be portable and may be carried on one or more predetermined routes to assess predefined points of measurement. In these many examples, the operating system that may be included in the MRDS 5082 may be Windows™, Linux™, or MacOS™ operating systems, or other similar operating systems. Further, in these arrangements, the operating system, modules for the operating system, and other needed libraries, data storage, and the like may be accessible wholly or partially through access to the cloud network facility 5080. In example embodiments, the MRDS 5082 may reside directly on the DAQ instrument 5002, especially in on-line system examples. In example embodiments, the DAQ instrument 5002 may be linked on an intra-network in a facility but may otherwise be behind a firewall. In further examples, the DAQ instrument 5002 may be linked to the cloud network facility 5080. In the various embodiments, one of the computers or mobile computing devices may be effectively designated the MRDS 5082 to which all of the other computing devices may feed it data such as one of the MRDS 6104, as depicted in Figures 31 and 32. In the many examples where the DAQ instrument 5002 may be deployed and configured to receive stream data in a swarm environment, one or more of the DAQ instruments 5002 may be effectively designated the MRDS 5082 to which all of the other computing devices may feed it data. In the many examples where the DAQ instrument 5002 may be deployed and configured to receive stream data in an environment where the methods and systems disclosed herein are intelligently assigning, controlling, adjusting, and re-adjusting data pools, computing resources, network bandwidth for local data collection, and the like, one or more of the DAQ instruments 5002 may be effectively designated the MRDS 5082 to which all of the other computing devices may feed it data.

[0446] With further reference to Figure 22, new raw streaming data, data that have been through extract, process, and align processes (EP data), and the like may be uploaded to one or more master raw data servers as needed or as scaled in various environments. In example embodiments, a master raw data server ("MRDS") 5700 may connect to and receive data from other master raw data servers such as the MRDS 5082. The MRDS 5700 may include a data distribution manager module 5702. In example embodiments, the new raw streaming data may be stored in the new stream data repository 5704. In many instances, like raw data streams stored on the DAQ instrument 5002, the new stream data repository 5704 and new extract and process data repository 5708 may be similarly configured as a temporary storage area.

[0447] In example embodiments, the MRDS 5700 may include a stream data analyzer module with an extract and process alignment module 5710. The analyzer module 5710 may be shown to be a more robust data analyzer and extractor than may be typically found on portable streaming DAQ instruments although it may be deployed on the DAQ instrument 5002 as well. In example embodiments, the analyzer module 5710 takes streaming data and instantiates it at a specific sampling rate and resolution similar to the local data control module 5062 on the DAQ instrument 5002. The specific sampling rate and resolution of the analyzer module 5710 may be based on

either user input 5712 or automated extractions from a multimedia probe ("MMP") and the probe control, sequence and analytical ("PCSA") information store 5714 and/or an identification mapping table 5718, which may require the user input 5712 if there is incomplete information regarding various forms of legacy data similar to as was detailed with the DAQ instrument 5002. In example embodiments, legacy data may be processed with the analyzer module 5710 and may be stored in one or more temporary holding areas such as a new legacy data repository 5720. One or more temporary areas may be configured to hold data until it is copied to an archive and verified. The analyzer 5710 module may also facilitate in-depth analysis by providing many varying types of signal processing tools including but not limited to filtering, Fourier transforms, weighting, resampling, envelope demodulation, wavelets, two-channel analysis, and the like. From this analysis, many different types of plots and mini-reports may be generated from a reports and plots module 5724. In example embodiments, data is sent to the processing, analysis, reports, and archiving ("PARA") server 5730 upon user initiation or in an automated fashion especially for on-line systems.

[0448] In example embodiments, a PARA server 5750 may connect to and receive data from other PARA servers such as the PARA server 5730. With reference to Figure 24, the PARA server 5730 may provide data to a supervisory module 5752 on the PARA server 5750 that may be configured to provide at least one of processing, analysis, reporting, archiving, supervisory, and similar functionalities. The supervisory module 5752 may also contain extract, process align functionality and the like. In example embodiments, incoming streaming data may first be stored in a raw data stream archive 5760 after being properly validated. Based on the analytical requirements derived from a multimedia probe ("MMP") and probe control, sequence and analytical ("PCSA") information store 5762 as well as user settings, data may be extracted, analyzed, and stored in an extract and process ("EP") raw data archive 5764. In example embodiments, various reports from a reports module 5768 are generated from the supervisory module 5752. The various reports from the reports module 5768 include trend plots of various smart bands, overalls along with statistical patterns, and the like. In example embodiments, the reports module 5768 may also be configured to compare incoming data to historical data. By way of these examples, the reports module 5768 may search for and analyze adverse trends, sudden changes, machinery defect patterns, and the like. In example embodiments, the PARA server 5750 may include an expert analysis module 5770 from which reports are generated and analysis may be conducted. Upon completion, archived data may be fed to a local master server ("LMS") 5772 via a server module 5774 that may connect to the local area network. In example embodiments, archived data may also be fed to the LMS 5772 via a cloud data management server ("CDMS") 5778 through a server module for a cloud network facility 5080. In example embodiments, the supervisory module 5752 on the PARA server 5750 may be configured to provide at least one of processing, analysis, reporting, archiving, supervisory, and similar functionalities from which alarms may be generated, rated, stored, modified, reassigned, and the like with an alarm generator module 5782.

[0449] Figure 24 depicts various embodiments that include a PARA server 5800 and its connection to LAN 5802. In example embodiments, one or more DAQ instruments such as the DAQ

instrument 5002 may receive and process analog data from one or more analog sensors 5710 that may be fed into the DAQ instrument 5002. As discussed herein, the DAQ instrument 5002 may create a digital stream of data based on the ingested analog data from the one or more analog sensors. The digital stream from the DAQ instrument 5002 may be uploaded to the MRDS 5082 and from there, it may be sent to the PARA server 5800 where multiple terminals, such as terminal 5810 5812, 5814, may each interface with it or the MRDS 5082 and view the data and/or analysis reports. In example embodiments, the PARA server 5800 may communicate with a network data server 5820 that may include a LMS 5822. In these examples, the LMS 5822 may be configured as an optional storage area for archived data. The LMS 5822 may also be configured as an external driver that may be connected to a PC or other computing device that may run the LMS 5822; or the LMS 5822 may be directly run by the PARA server 5800 where the LMS 5822 may be configured to operate and coexist with the PARA server 5800. The LMS 5822 may connect with a raw data stream archive 5824, an extract and process ("EP") raw data archive 5828, and a MMP and probe control, sequence and analytical ("PCSA") information store 5830. In example embodiments, a CDMS 5832 may also connect to the LAN 5802 and may also support the archiving of data.

[0450] In example embodiments, portable connected devices 5850 such as a tablet 5852 and a smart phone 5854 may connect the CDMS 5832 using web APIs 5860 and 5862, respectively, as depicted in Figure 25. The APIs 5860, 5862 may be configured to execute in a browser and may permit access via a cloud network facility 5870 of all (or some of) the functions previously discussed as accessible through the PARA Server 5800. In example embodiments, computing devices of a user 5880 such as computing devices 5882, 5884, 5888 may also access the cloud network facility 5870 via a browser or other connection in order to receive the same functionality. In example embodiments, thin-client apps which do not require any other device drivers and may be facilitated by web services supported by cloud services 5890 and cloud data 5892. In many examples, the thin-client apps may be developed and reconfigured using, for example, the visual high-level LabVIEW™ programming language with NXG™ Web-based virtual interface subroutines. In example embodiments, thin client apps may provide high-level graphing functions such as those supported by LabVIEW™ tools. In example embodiments, the LabVIEW™ tools may generate JSCRIPT™ code and JAVA™ code that may be edited post-compilation. The NXG™ tools may generate Web VI's that may not require any specialized driver and only some RESTful™ services which may be readily installed from any browser. It will be appreciated in light of the disclosure that because various applications may be run inside a browser, the applications may be run on any operating system, such as Windows™, Linux™, and Android™ operating systems especially for personal devices, mobile devices, portable connected devices, and the like.

[0451] In example embodiments, the CDMS 5832 is depicted in greater detail in Figure 26. In example embodiments, the CDMS 5832 may provide all of the data storage and services that the PARA Server 5800 (Figure 34) may provide. In contrast, all of the API's may be web API's which may run in a browser and all other apps may run on the PARA Server 5800 or the DAQ instrument

5002 which may typically be Windows™, Linux™ or other similar operating systems. In example embodiments, the CDMS 5832 includes at least one of or combinations of the following functions: the CDMS 5832 may include a cloud GUI 5900 that may be configured to provide access to all data plots including trend, waveform, spectra, envelope, transfer function, logs of measurement events, analysis including expert, utilities, and the like. In example embodiments, the CDMS 5832 may include a cloud data exchange 5902 configured to facilitate the transfer of data to and from the cloud network facility 5870. In example embodiments, the CDMS 5832 may include a cloud plots/trends module 5904 that may be configured to show all plots via web apps including trend, waveform, spectra, envelope, transfer function, and the like. In example embodiments, the CDMS 5832 may include a cloud reporter 5908 that may be configured to provide all analysis reports, logs, expert analysis, trend plots, statistical information, and the like. In example embodiments, the CDMS 5832 may include a cloud alarm module 5910. Alarms from the cloud alarm module 5910 may be generated and may be sent to various devices 5920 via email, texts, or other messaging mechanisms. From the various modules, data may be stored in new data 5914. The various devices 5920 may include a terminal 5922, portable connected device 5924, or a tablet 5928. The alarms from the cloud alarm module are designed to be interactive so that the end user may acknowledge alarms in order to avoid receiving redundant alarms and also to see significant context-sensitive data from the alarm points that may include spectra, waveform statistical info, and the like.

[0452] In example embodiments, a relational database server ("RDS") 5930 may be used to access all of the information from a MMP and PCSA information store 5932. As with the PARA server 5800 (Figure 26), information from the information store 5932 may be used with an EP and align module 5934, a data exchange 5938 and the expert system 5940. In example embodiments, a raw data stream archive 5942 and extract and process raw data archive 5944 may also be used by the EP align 5934, the data exchange 5938 and the expert system 5940 as with the PARA server 5800. In example embodiments, new stream raw data 5950, new extract and process raw data 5952, and new data 5954 (essentially all other raw data such as overalls, smart bands, stats, and data from the information store 5932) are directed by the CDMS 5832.

[0453] In example embodiments, the streaming data may be linked with the RDS 5930 and the MMP and PCSA information store 5932 using a technical data management streaming ("TDMS") file format. In example embodiments, the information store 5932 may include tables for recording at least portions of all measurement events. By way of these examples, a measurement event may be any single data capture, a stream, a snapshot, an averaged level, or an overall level. Each of the measurement events in addition to point identification information may also have a date and time stamp. In example embodiments, a link may be made between the streaming data, the measurement event, and the tables in the information store 5932 using the TDMS format. By way of these examples, the link may be created by storing unique measurement point identification codes with a file structure having the TDMS format by including and assigning TDMS properties. In example embodiments, a file with the TDMS format may allow for three levels of hierarchy. By way of these examples, the three levels of hierarchy may be root, group, and channel. It will be appreciated

in light of the disclosure that the Mimosa™ database schema may be, in theory, unlimited. With that said, there are advantages to limited TDMS hierarchies. In the many examples, the following properties may be proposed for adding to the TDMS Stream structure while using a Mimosa Compatible database schema.

[0454] Root Level: Global ID 1: Text String (This could be a unique ID obtained from the web.); Global ID 2: Text String (This could be an additional ID obtained from the web.); Company Name: Text String; Company ID: Text String; Company Segment ID: 4-byte Integer; Company Segment ID: 4-byte Integer; Site Name: Text String; Site Segment ID: 4-byte Integer; Site Asset ID: 4-byte Integer; Route Name: Text String; Version Number: Text String

[0455] Group Level: Section 1 Name: Text String; Section 1 Segment ID: 4-byte Integer; Section 1 Asset ID: 4-byte Integer; Section 2 Name: Text String; Section 2 Segment ID: 4-byte Integer; Section 2 Asset ID: 4-byte Integer; Machine Name: Text String; Machine Segment ID: 4-byte Integer; Machine Asset ID: 4-byte Integer; Equipment Name: Text String; Equipment Segment ID: 4-byte Integer; Equipment Asset ID: 4-byte Integer; Shaft Name: Text String; Shaft Segment ID: 4-byte Integer; Shaft Asset ID: 4-byte Integer; Bearing Name: Text String; Bearing Segment ID: 4-byte Integer; Bearing Asset ID: 4-byte Integer; Probe Name: Text String; Probe Segment ID: 4-byte Integer; Probe Asset ID: 4-byte Integer

[0456] Channel Level: Channel #: 4-byte Integer; Direction: 4-byte Integer (in certain examples may be text); Data Type: 4-byte Integer; Reserved Name 1: Text String; Reserved Segment ID 1: 4-byte Integer; Reserved Name 2: Text String; Reserved Segment ID 2: 4-byte Integer; Reserved Name 3: Text String; Reserved Segment ID 3: 4-byte Integer

[0457] In example embodiments, the file with the TDMS format may automatically use property or asset information and may make an index file out of the specific property and asset information to facilitate database searches, may offer a compromise for storing voluminous streams of data because it may be optimized for storing binary streams of data but may also include some minimal database structure making many standard SQL operations feasible, but the TDMS format and functionality discussed herein may not be as efficient as a full-fledged SQL relational database. The TDMS format, however, may take advantage of both worlds in that it may balance between the class or format of writing and storing large streams of binary data efficiently and the class or format of a fully relational database, which facilitates searching, sorting and data retrieval. In example embodiments, an optimum solution may be found in that metadata required for analytical purposes and extracting prescribed lists with panel conditions for stream collection may be stored in the RDS 5930 by establishing a link between the two database methodologies. By way of these examples, relatively large analog data streams may be stored predominantly as binary storage in the raw data stream archive 5942 for rapid stream loading but with inherent relational SQL type hooks, formats, conventions, or the like. The files with the TDMS format may also be configured to incorporate DIAdem™ reporting capability of LabVIEW™ software in order to provide a further mechanism to conveniently and rapidly facilitate accessing the analog or the streaming data.

[0458] The methods and systems disclosed herein may include, connect to, or be integrated with a virtual data acquisition instrument and in the many embodiments, Figure 27 shows methods and systems that include a virtual streaming DAQ instrument 6000 also known as a virtual DAQ instrument, a VRDS, or a VSDAQ. In contrast to the DAQ instrument 5002 (Figure 18), the virtual DAQ instrument 6000 may be configured so to only include one native application. In the many examples, the one permitted and one native application may be the DAQ driver module 6002 that may manage all communications with the DAQ Device 6004 which may include streaming capabilities. In example embodiments, other applications, if any, may be configured as thin client web applications such as RESTful™ web services. The one native application, or other applications or services, may be accessible through the DAQ Web API 6010. The DAQ Web API 6010 may run in or be accessible through various web browsers.

[0459] In example embodiments, storage of streaming data, as well as the extraction and processing of streaming data into extract and process data, may be handled primarily by the DAQ driver services 6012 under the direction of the DAQ Web API 6010. In example embodiments, the output from sensors of various types including vibration, temperature, pressure, ultrasound and so on may be fed into the instrument inputs of the DAQ device 6004. In example embodiments, the signals from the output sensors may be signal conditioned with respect to scaling and filtering and digitized with an analog to a digital converter. In example embodiments, the signals from the output sensors may be signals from all relevant channels simultaneously sampled at a rate sufficient to perform the maximum desired frequency analysis. In example embodiments, the signals from the output sensors may be sampled for a relatively long time, gap-free, as one continuous stream so as to enable a wide array of further post-processing at lower sampling rates with sufficient samples. In further examples, streaming frequency may be adjusted (and readjusted) to record streaming data at non-evenly spaced recording. For temperature data, pressure data, and other similar data that may be relatively slow, varying delta times between samples may further improve quality of the data. By way of the above examples, data may be streamed from a collection of points and then the next set of data may be collected from additional points according to a prescribed sequence, route, path, or the like. In the many examples, the portable sensors may be moved to the next location according to the prescribed sequence but not necessarily all of them as some may be used for reference phase or otherwise. In further examples, a multiplexer 6020 may be used to switch to the next collection of points or a mixture of the two methods may be combined.

[0460] In example embodiments, the sequence and panel conditions that may be used to govern the data collection process using the virtual DAQ instrument 6000 may be obtained from the MMP PCSA information store 6022. The MMP PCSA information store 6022 may include such items as the hierarchical structural relationships of the machine, i.e., a machine contains pieces of equipment in which each piece of equipment contains shafts and each shaft is associated with bearings, which may be monitored by specific types of transducers or probes according to a specific prescribed sequence (routes, path, etc.) with specific panel conditions. By way of these examples, the panel conditions may include hardware specific switch settings or other collection parameters such as sampling rate, AC/DC coupling, voltage range and gain, integration, high and low pass

filtering, anti-aliasing filtering, ICP™ transducers and other integrated-circuit piezoelectric transducers, 4-20 mA loop sensors, and the like. The information store 6022 includes other information that may be stored in what would be machinery specific features that would be important for proper analysis including the number of gear teeth for a gear, the number of blades in a pump impeller, the number of motor rotor bars, bearing specific parameters necessary for calculating bearing frequencies, 1x rotating speed (RPMs) of all rotating elements, and the like.

[0461] Upon direction of the DAQ Web API 6010 software, digitized waveforms may be uploaded using the DAQ driver services 6012 of the virtual DAQ instrument 6000. In example embodiments, data may then be fed into an RLN data and control server 6030 that may store the stream data into a network stream data repository 6032. Unlike the DAQ instrument 5002, the server 6030 may run from within the DAQ driver module 6002. It will be appreciated in light of the disclosure that a separate application may require drivers for running in the native operating system and for this instrument only the instrument driver may run natively. In many examples, all other applications may be configured to be browser based. As such, a relevant network variable may be very similar to a LabVIEW™ shared or network stream variable which may be designed to be accessed over one or more networks or via web applications.

[0462] In example embodiments, the DAQ web API 6010 may also direct the local data control application 6034 to extract and process the recently obtained streaming data and, in turn, convert it to the same or lower sampling rates of sufficient length to provide the desired resolution. This data may be converted to spectra, then averaged and processed in a variety of ways and stored as EP data, such as on an EP data repository 6040. The EP data repository 6040 may, in certain example embodiments, only be meant for temporary storage. It will be appreciated in light of the disclosure that legacy data may require its own sampling rates and resolution and often this sampling rate may not be integer proportional to the acquired sampling rate especially for order-sampled data whose sampling frequency is related directly to an external frequency. The external frequency may typically be the running speed of the machine or its internal componentry, rather than the more-standard sampling rates produced by the internal crystals, clock functions, and the like of the (e.g., values of Fmax of 100, 200, 500, 1K, 2K, 5K, 10K, 20K and so on) of the DAQ instrument 5002, 6000. In example embodiments, the EP align component of the local data control application 6034 is able to fractionally adjust the sampling rate to the non-integer ratio rates that may be more applicable to legacy data sets and therefore drive compatibility with legacy systems. In example embodiments, the fractional rates may be converted to integer ratio rates more readily because the length of the data to be processed (or at least that portion of the greater stream of data) is adjustable because of the depth and content of the original acquired streaming data by the DAQ instrument 5002, 6000. It will be appreciated in light of the disclosure that if the data was not streamed and just stored as traditional snap-shots of spectra with the standard values of Fmax, it may very well be impossible to retroactively and accurately convert the acquired data to the order-sampled data. In example embodiments, the stream data may be converted, especially for legacy data purposes, to the proper sampling rate and resolution as described and stored in the EP legacy data repository 6042. To support legacy data identification scenarios, a user input 6044 may be

included if there is no automated process for identification translation. In example embodiments, one such automated process for identification translation may include importation of data from a legacy system that may contain a fully standardized format such as the Mimosa™ format and sufficient identification information to complete an ID Mapping Table 6048. In further examples, the end user, a legacy data vendor, a legacy data storage facility, or the like may be able to supply enough info to complete (or sufficiently complete) relevant portions of the ID Mapping Table 6048 to provide, in turn, the database schema for the raw data of the legacy system so it may be readily ingested, saved, and used for analytics in the current systems disclosed herein.

[0463] Figure 28 depicts further embodiments and details of the virtual DAQ Instrument 6000. In these examples, the DAQ Web API 6010 may control the data collection process as well as its sequence. The DAQ Web API 6010 may provide the capability for editing this process, viewing plots of the data, controlling the processing of that data and viewing the output in all its myriad forms, analyzing the data, including the expert analysis, communicating with external devices via the DAQ driver module 6002, as well as communicating with and transferring both streaming data and EP data to one or more cloud network facilities 5080 whenever possible. In example embodiments, the virtual DAQ instrument itself and the DAQ Web API 6010 may run independently of access to cloud network facilities 5080 when local demands may require or simply as a result of there being no outside connectivity such use throughout a proprietary industrial setting that prevents such signals. In example embodiments, the DAQ Web API 6010 may also govern the movement of data, its filtering, as well as many other housekeeping functions.

[0464] The virtual DAQ Instrument 6000 may also include an expert analysis module 6052. In example embodiments, the expert analysis module 6052 may be a web application or other suitable module that may generate reports 6054 that may use machine or measurement point specific information from the MMP PCSA information store 6022 to analyze stream data 6058 using the stream data analyzer module 6050. In example embodiments, supervisory control of the module 6052 may be provided by the DAQ Web API 6010. In example embodiments, the expert analysis may also be supplied (or supplemented) via the expert system module 5940 that may be resident on one or more cloud network facilities that are accessible via the CDMS 5832. In many examples, expert analysis via the cloud may be preferred over local systems such as expert analysis module 6052 for various reasons, such as the availability and use of the most up-to-date software version, more processing capability, a bigger volume of historical data to reference and the like. It will be appreciated in light of the disclosure that it may be important to offer expert analysis when an internet connection cannot be established so as to provide a redundancy, when needed, for seamless and time efficient operation. In example embodiments, this redundancy may be extended to all of the discussed modular software applications and databases where applicable so each module discussed herein may be configured to provide redundancy to continue operation in the absence of an internet connection.

[0465] Figure 29 depicts further embodiments and details of many virtual DAQ instruments existing in an online system and connecting through network endpoints through a central DAQ instrument to one or more cloud network facilities. In example embodiments, a master DAQ

instrument with network endpoint 6060 is provided along with additional DAQ instruments such as a DAQ instrument with network endpoint 6062, a DAQ instrument with network endpoint 6064, and a DAQ instrument with network endpoint 6068. The master DAQ instrument with network endpoint 6060 may connect with the other DAQ instruments with network endpoints 6062, 6064, 6068 over LAN 6070. It will be appreciated that each of the instruments 6060, 6062, 6064, 6068 may include personal computer, a connected device, or the like that include Windows™, Linux™, or other suitable operating systems to facilitate ease of connection of devices utilizing many wired and wireless network options such as Ethernet, wireless 802.11g, 900 MHz wireless (e.g., for better penetration of walls, enclosures and other structural barriers commonly encountered in an industrial setting), as well as a myriad of other things permitted by the use of off-the-shelf communication hardware when needed.

[0466] Figure 30 depicts further embodiments and details of many functional components of an endpoint that may be used in the various settings, environments, and network connectivity settings. The endpoint includes endpoint hardware modules 6080. In example embodiments, the endpoint hardware modules 6080 may include one or more multiplexers 6082, a DAQ instrument 6084, as well as a computer 6088, computing device, PC, or the like that may include the multiplexers, DAQ instruments, and computers, connected devices and the like, as disclosed herein. The endpoint software modules 6090 include a data collector application (DCA) 6092 and a raw data server (RDS) 6094. In example embodiments, DCA 6092 may be similar to the DAQ API 5052 (Figure 18) and may be configured to be responsible for obtaining stream data from the DAQ device 6084 and storing it locally according to a prescribed sequence or upon user directives. In the many examples, the prescribed sequence or user directives may be a LabVIEW™ software app that may control and read data from the DAQ instruments. For cloud based online systems, the stored data in many embodiments may be network accessible. In many examples, LabVIEW™ tools may be used to accomplish this with a shared variable or network stream (or subsets of shared variables). Shared variables and the affiliated network streams may be network objects that may be optimized for sharing data over the network. In many embodiments, the DCA 6092 may be configured with a graphic user interface that may be configured to collect data as efficiently and fast as possible and push it to the shared variable and its affiliated network stream. In example embodiments, the endpoint raw data server 6094 may be configured to read raw data from the single-process shared variable and may place it with a master network stream. In example embodiments, a raw stream of data from portable systems may be stored locally and temporarily until the raw stream of data is pushed to the MRDS 5082 (Figure 18). It will be appreciated in light of the disclosure that on-line system instruments on a network can be termed endpoints whether local or remote or associated with a local area network or a wide area network. For portable data collector applications that may or may not be wirelessly connected to one or more cloud network facilities, the endpoint term may be omitted as described so as to detail an instrument that may not require network connectivity.

[0467] Figure 31 depicts further embodiments and details of multiple endpoints with their respective software blocks with at least one of the devices configured as master blocks. Each of

the blocks may include a data collector application ("DCA") 7000 and a raw data server ("RDS") 7002. In example embodiments, each of the blocks may also include a master raw data server module ("MRDS") 7004, a master data collection and analysis module ("MDCA") 7008, and a supervisory and control interface module ("SCI") 7010. The MRDS 7004 may be configured to read network stream data (at a minimum) from the other endpoints and may forward it up to one or more cloud network facilities via the CDMS 5832 including the cloud services 5890 and the cloud data 5892. In example embodiments, the CDMS 5832 may be configured to store the data and to provide web, data, and processing services. In these examples, this may be implemented with a LabVIEW™ application that may be configured to read data from the network streams or share variables from all of the local endpoints, write them to the local host PC, local computing device, connected device, or the like, as both a network stream and file with TDMS™ formatting. In example embodiments, the CDMS 5832 may also be configured to then post this data to the appropriate buckets using the LabVIEW or similar software that may be supported by S3™ web service from the Amazon Web Services ("AWS™") on the Amazon™ web server, or the like and may effectively serve as a back-end server. In the many examples, different criteria may be enabled or may be set up for when to post data, create or adjust schedules, create or adjust event triggering including a new data event, create a buffer full message, create or more alarms messages, and the like.

[0468] In example embodiments, the MDCA 7008 may be configured to provide automated as well as user-directed analyses of the raw data that may include tracking and annotating specific occurrence and in doing so, noting where reports may be generated and alarms may be noted. In example embodiments, the SCI 7010 may be an application configured to provide remote control of the system from the cloud as well as the ability to generate status and alarms. In example embodiments, the SCI 7010 may be configured to connect to, interface with, or be integrated into a supervisory control and data acquisition ("SCADA") control system. In example embodiments, the SCI 7010 may be configured as a LabVIEW™ application that may provide remote control and status alerts that may be provided to any remote device that may connect to one or more of the cloud network facilities 5870.

[0469] In example embodiments, the equipment that is being monitored may include RFID tags that may provide vital machinery analysis background information. The RFID tags may be associated with the entire machine or associated with the individual componentry and may be substituted when certain parts of the machine are replaced, repaired, or rebuilt. The RFID tags may provide permanent information relevant to the lifetime of the unit or may also be re-flashed to update with at least a portion of new information. In many embodiments, the DAQ instruments 5002 disclosed herein may interrogate the one or more RFID chips to learn of the machine, its componentry, its service history, and the hierarchical structure of how everything is connected including drive diagrams, wire diagrams, and hydraulic layouts. In example embodiments, some of the information that may be retrieved from the RFID tags includes manufacturer, machinery type, model, serial number, model number, manufacturing date, installation date, lot numbers, and the like. By way of these examples, machinery type may include the use of a Mimosa™ format

table including information about one or more of the following motors, gearboxes, fans, and compressors. The machinery type may also include the number of bearings, their type, their positioning, and their identification numbers. The information relevant to one or more fans includes fan type, number of blades, number of vanes, and number of belts. It will be appreciated in light of the disclosure that other machines and their componentry may be similarly arranged hierarchically with relevant information all of which may be available through interrogation of one or more RFID chips associated with the one or more machines.

[0470] In example embodiments, data collection in an industrial environment may include routing analog signals from a plurality of sources, such as analog sensors, to a plurality of analog signal processing circuits. Routing of analog signals may be accomplished by an analog crosspoint switch that may route any of a plurality of analog input signals to any of a plurality of outputs, such as to analog and/or digital outputs. Routing of inputs to outputs in an analog signal crosspoint switch in an industrial environment may be configurable, such as by an electronic signal to which a switch portion of the analog crosspoint switch is responsive.

[0471] In example embodiments, the analog crosspoint switch may receive analog signals from a plurality of analog signal sources in the industrial environment. Analog signal sources may include sensors that produce an analog signal. Sensors that produce an analog signal that may be switched by the analog crosspoint switch may include sensors that detect a condition and convert it to an analog signal that may be representative of the condition, such as converting a condition to a corresponding voltage. Exemplary conditions that may be represented by a variable voltage may include temperature, friction, sound, light, torque, revolutions-per-minute, mechanical resistance, pressure, flow rate, and the like, including any of the conditions represented by inputs sources and sensors disclosed throughout this disclosure and the documents incorporated herein by reference. Other forms of analog signal may include electrical signals, such as variable voltage, variable current, variable resistance, and the like.

[0472] In example embodiments, the analog crosspoint switch may preserve one or more aspects of an analog signal being input to it in an industrial environment. Analog circuits integrated into the switch may provide buffered outputs. The analog circuits of the analog crosspoint switch may follow an input signal, such as an input voltage to produce a buffered representation on an output. This may alternatively be accomplished by relays (mechanical, solid state, and the like) that allow an analog voltage or current present on an input to propagate to a selected output of the analog switch.

[0473] In example embodiments, an analog crosspoint switch in an industrial environment may be configured to switch any of a plurality of analog inputs to any of a plurality of analog outputs. An example embodiment includes a MIMO, multiplexed configuration. An analog crosspoint switch may be dynamically configurable so that changes to the configuration causes a change in the mapping of inputs to outputs. A configuration change may apply to one or more mappings so that a change in mapping may result in one or more of the outputs being mapped to different input than before the configuration change.

[0474] In example embodiments, the analog crosspoint switch may have more inputs than outputs, so that only a subset of inputs can be routed to outputs concurrently. In other embodiments, the analog crosspoint switch may have more outputs than inputs, so that either a single input may be made available currently on multiple outputs, or at least one output may not be mapped to any input.

[0475] In example embodiments, an analog crosspoint switch in an industrial environment may be configured to switch any of a plurality of analog inputs to any of a plurality of digital outputs. To accomplish conversion from analog inputs to digital outputs, an analog-to-digital converter circuit may be configured on each input, each output, or at intermediate points between the input(s) and output(s) of the analog crosspoint switch. Benefits of including digitization of analog signals in an analog crosspoint switch that may be located close to analog signal sources may include reducing signal transport costs and complexity that digital signal communication has over analog, reducing energy consumption, facilitating detection and regulation of aberrant conditions before they propagate throughout an industrial environment, and the like. Capturing analog signals close to their source may also facilitate improved signal routing management that is more tolerant of real world effects such as requiring that multiple signals be routed simultaneously. In this example, a portion of the signals can be captured (and stored) locally while another portion can be transferred through the data collection network. Once the data collection network has available bandwidth, the locally stored signals can be delivered, such as with a time stamp indicating the time at which the data was collected. This technique may be useful for applications that have concurrent demand for data collection channels that exceed the number of channels available. Sampling control may also be based on an indication of data worth sampling. As an example, a signal source, such as a sensor in an industrial environment may provide a data valid signal that transmits an indication of when data from the sensor is available.

[0476] In example embodiments, mapping inputs of the analog crosspoint switch to outputs may be based on a signal route plan for a portion of the industrial environment that may be presented to the crosspoint switch. The signal route plan may be used in a method of data collection in the industrial environment that may include routing a plurality of analog signals along a plurality of analog signal paths. The method may include connecting the plurality of analog signals individually to inputs of the analog crosspoint switch that may be configured with a route plan. The crosspoint switch may, responsively to the configured route plan, route a portion of the plurality of analog signals to a portion of the plurality of analog signal paths.

[0477] In example embodiments, the analog crosspoint switch may include at least one high current output drive circuit that may be suitable for routing the analog signal along a path that requires high current. In example embodiments, the analog crosspoint switch may include at least one voltage-limited input that may facilitate protecting the analog crosspoint switch from damage due to excessive analog input signal voltage. In example embodiments, the analog crosspoint switch may include at least one current limited input that may facilitate protecting the analog crosspoint switch from damage due to excessive analog input current. The analog crosspoint switch

may comprise a plurality of interconnected relays that may facilitate routing the input(s) to the output(s) with little or no substantive signal loss.

[0478] In example embodiments, an analog crosspoint switch may include processing functionality, such as signal processing and the like (e.g., a programmed processor, special purpose processor, a digital signal processor, and the like) that may detect one or more analog input signal conditions. In response to such detection, one or more actions may be performed, such as setting an alarm, sending an alarm signal to another device in the industrial environment, changing the crosspoint switch configuration, disabling one or more outputs, powering on or off a portion of the switch, changing a state of an output, such as a general purpose digital or analog output, and the like. In example embodiments, the switch may be configured to process inputs for producing a signal on one or more of the outputs. The inputs to use, processing algorithm for the inputs, condition for producing the signal, output to use, and the like may be configured in a data collection template.

[0479] In example embodiments, an analog crosspoint switch may comprise greater than 32 inputs and greater than 32 outputs. A plurality of analog crosspoint switches may be configured so that even though each switch offers fewer than 32 inputs and 32 outputs it may be configured to facilitate switching any of 32 inputs to any of 32 outputs spread across the plurality of crosspoint switches.

[0480] In example embodiments, an analog crosspoint switch suitable for use in an industrial environment may comprise four or fewer inputs and four or fewer outputs. Each output may be configurable to produce an analog output that corresponds to the mapped analog input or it may be configured to produce a digital representation of the corresponding mapped input.

[0481] In example embodiments, an analog crosspoint switch for use in an industrial environment may be configured with circuits that facilitate replicating at least a portion of attributes of the input signal, such as current, voltage range, offset, frequency, duty cycle, ramp rate, and the like while buffering (e.g., isolating) the input signal from the output signal. Alternatively, an analog crosspoint switch may be configured with unbuffered inputs/outputs, thereby effectively producing a bi-directional based crosspoint switch).

[0482] In example embodiments, an analog crosspoint switch for use in an industrial environment may include protected inputs that may be protected from damaging conditions, such as through use of signal conditioning circuits. Protected inputs may prevent damage to the switch and to downstream devices to which the switch outputs connect. As an example, inputs to such an analog crosspoint switch may include voltage clipping circuits that prevent a voltage of an input signal from exceeding an input protection threshold. An active voltage adjustment circuit may scale an input signal by reducing it uniformly so that a maximum voltage present on the input does not exceed a safe threshold value. As another example, inputs to such an analog crosspoint switch may include current shunting circuits that cause current beyond a maximum input protection current threshold to be diverted through protection circuits rather than enter the switch. Analog switch inputs may be protected from electrostatic discharge and/or lightning strikes. Other signal conditioning functions that may be applied to inputs to an analog crosspoint switch may include

voltage scaling circuitry that attempts to facilitate distinguishing between valid input signals and low voltage noise that may be present on the input. However, in example embodiments, inputs to the analog crosspoint switch may be unbuffered and/or unprotected to make the least impact on the signal. Signals such as alarm signals, or signals that cannot readily tolerate protection schemes, such as those schemes described herein may be connected to unbuffered inputs of the analog crosspoint switch.

[0483] In example embodiments, an analog crosspoint switch may be configured with circuitry, logic, and/or processing elements that may facilitate input signal alarm monitoring. Such an analog crosspoint switch may detect inputs meeting alarm conditions and in response thereto, switch inputs, switch mapping of inputs to outputs, disable inputs, disable outputs, issue an alarm signal, activate/deactivate a general-purpose output, or the like.

[0484] In example embodiments, a system for collecting data in an industrial environment may include an analog crosspoint switch that may be adapted to selectively power up or down portions of the analog crosspoint switch or circuitry associated with the analog crosspoint switch, such as input protection devices, input conditioning devices, switch control devices and the like. Portions of the analog crosspoint switch that may be powered on/off may include outputs, inputs, sections of the switch and the like. In an example, an analog crosspoint switch may include a modular structure that may separate portions of the switch into independently powered sections. Based on conditions, such as an input signal meeting a criterion or a configuration value being presented to the analog crosspoint switch, one or more modular sections may be powered on / off.

[0485] In example embodiments, a system for collecting data in an industrial environment may include an analog crosspoint switch that may be adapted to perform signal processing including, without limitation, providing a voltage reference for detecting an input crossing the voltage reference (e.g., zero volts for detecting zero-crossing signals), a phase-lock loop to facilitate capturing slow frequency signals (e.g., low-speed revolution-per-minute signals and detecting their corresponding phase), deriving input signal phase relative to other inputs, deriving input signal phase relative to a reference (e.g., a reference clock), deriving input signal phase relative to detected alarm input conditions and the like. Other signal processing functions of such an analog crosspoint switch may include oversampling of inputs for delta-sigma A/D, to produce lower sampling rate outputs, to minimize AA filter requirements and the like. Such an analog crosspoint switch may support long block sampling at a constant sampling rate even as inputs are switched, which may facilitate input signal rate independence and reduce complexity of sampling scheme(s). A constant sampling rate may be selected from a plurality of rates that may be produced by a circuit, such as a clock divider circuit that may make available a plurality of components of a reference clock.

[0486] In example embodiments, a system for collecting data in an industrial environment may include an analog crosspoint switch that may be adapted to support implementing data collection / data routing templates in the industrial environment. The analog crosspoint switch may implement a data collection / data routing template based on conditions in the industrial environment that it may detect or derive, such as an input signal meeting one or more criteria (e.g., transition of a

signal from a first condition to a second, lack of transition of an input signal within a predefined time interface (e.g., inactive input) and the like).

[0487] In example embodiments, a system for collecting data in an industrial environment may include an analog crosspoint switch that may be adapted to be configured from a portion of a data collection template. Configuration may be done automatically (without needing human intervention to perform a configuration action or change in configuration), such as based on a time parameter in the template and the like. Configuration may be done remotely, e.g., by sending a signal from a remote location that is detectable by a switch configuration feature of the analog crosspoint switch. Configuration may be done dynamically, such as based on a condition that is detectable by a configuration feature of the analog crosspoint switch (e.g., a timer, an input condition, an output condition, and the like). In example embodiments, information for configuring an analog crosspoint switch may be provided in a stream, as a set of control lines, as a data file, as an indexed data set, and the like. In example embodiments, configuration information in a data collection template for the switch may include a list of each input and a corresponding output, a list of each output function (active, inactive, analog, digital and the like), a condition for updating the configuration (e.g., an input signal meeting a condition, a trigger signal, a time (relative to another time/event/state, or absolute), a duration of the configuration, and the like. In example embodiments, configuration of the switch may be input signal protocol aware so that switching from a first input to a second input for a given output may occur based on the protocol. In an example, a configuration change may be initiated with the switch to switch from a first video signal to a second video signal. The configuration circuitry may detect the protocol of the input signal and switch to the second video signal during a synchronization phase of the video signal, such as during horizontal or vertical refresh. In other examples, switching may occur when one or more of the inputs are at zero volts. This may occur for a sinusoidal signal that transitions from below zero volts to above zero volts.

[0488] In example embodiments, a system for collecting data in an industrial environment may include an analog crosspoint switch that may be adapted to provide digital outputs by converting analog signals input to the switch into digital outputs. Converting may occur after switching the analog inputs based on a data collection template and the like. In example embodiments, a portion of the switch outputs may be digital and a portion may be analog. Each output, or groups thereof, may be configurable as analog or digital, such as based on analog crosspoint switch output configuration information included in or derived from a data collection template. Circuitry in the analog crosspoint switch may sense an input signal voltage range and intelligently configure an analog-to-digital conversion function accordingly. As an example, a first input may have a voltage range of 12 volts and a second input may have a voltage range of 24 volts. Analog-to-digital converter circuits for these inputs may be adjusted so that the full range of the digital value (e.g., 256 levels for an 8-bit signal) will map substantially linearly to 12 volts for the first input and 24 volts for the second input.

[0489] In example embodiments, an analog crosspoint switch may automatically configure input circuitry based on characteristics of a connected analog signal. Examples of circuitry configuration

may include setting a maximum voltage, a threshold based on a sensed maximum threshold, a voltage range above and/or below a ground reference, an offset reference, and the like. The analog crosspoint switch may also adapt inputs to support voltage signals, current signals, and the like. The analog crosspoint switch may detect a protocol of an input signal, such as a video signal protocol, audio signal protocol, digital signal protocol, protocol based on input signal frequency characteristics, and the like. Other aspects of inputs of the analog crosspoint switch that may be adapted based on the incoming signal may include a duration of sampling of the signal, and comparator or differential type signals, and the like.

[0490] In example embodiments, an analog crosspoint switch may be configured with functionality to counteract input signal drift and/or leakage that may occur when an analog signal is passed through it over a long period of time without changing value (e.g., a constant voltage). Techniques may include voltage boost, current injection, periodic zero referencing (e.g., temporarily connecting the input to a reference signal, such as ground, applying a high resistance pathway to the ground reference, and the like).

[0491] In example embodiments, a system for data collection in an industrial environment may include an analog crosspoint switch deployed in an assembly line comprising conveyers and/or lifters. A power roller conveyor system includes many rollers that deliver products along a path. There may be many points along the path that may be monitored for proper operation of the rollers, load being placed on the rollers, accumulation of products, and the like. A power roller conveyor system may also facilitate moving products through longer distances and therefore may have a large number of products in transport at once. A system for data collection in such an assembly environment may include sensors that detect a wide range of conditions as well as at a large number of positions along the transport path. As a product progresses down the path, some sensors may be active and others, such as those that the product has passed may be inactive. A data collection system may use an analog crosspoint switch to select only those sensors that are currently or anticipated to be active by switching from inputs that connect to inactive sensors to those that connect to active sensors and thereby provide the most useful sensor signals to data detection and/or collection and/or processing facilities. In example embodiments, the analog crosspoint switch may be configured by a conveyor control system that monitors product activity and instructs the analog crosspoint switch to direct different inputs to specific outputs based on a control program or data collection template associated with the assembly environment.

[0492] In example embodiments, a system for data collection in an industrial environment may include an analog crosspoint switch deployed in a factory comprising use of fans as industrial components. In example embodiments, fans in a factory setting may provide a range of functions including drying, exhaust management, clean air flow and the like. In an installation of a large number of fans, monitoring fan rotational speed, torque, and the like may be beneficial to detect an early indication of a potential problem with air flow being produced by the fans. However, concurrently monitoring each of these elements for a large number of fans may be inefficient. Therefore, sensors, such as tachometers, torque meters, and the like may be disposed at each fan and their analog output signal(s) may be provided to an analog crosspoint switch. With a limited

number of outputs, or at least a limited number of systems that can process the sensor data, the analog crosspoint switch may be used to select among the many sensors and pass along a subset of the available sensor signals to data collection, monitoring, and processing systems. In an example, sensor signals from sensors disposed at a group of fans may be selected to be switched onto crosspoint switch outputs. Upon satisfactory collection and/or processing of the sensor signals for this group of fans, the analog crosspoint switch may be reconfigured to switch signals from another group of fans to be processed.

[0493] In example embodiments, a system for data collection in an industrial environment may include an analog crosspoint switch deployed as an industrial component in a turbine-based power system. Monitoring for vibration in turbine systems, such as hydro-power systems, has been demonstrated to provide advantages in reduction in down time. However, with a large number of areas to monitor for vibration, particularly for on-line vibration monitoring, including relative shaft vibration, bearings absolute vibration, turbine cover vibration, thrust bearing axial vibration, stator core vibrations, stator bar vibrations, stator end winding vibrations, and the like, it may be beneficial to select among this list over time, such as taking samples from sensors for each of these types of vibration a few at a time. A data collection system that includes an analog crosspoint switch may provide this capability by connecting each vibration sensor to separate inputs of the analog crosspoint switch and configuring the switch to output a subset of its inputs. A vibration data processing system, such as a computer, may determine which sensors to pass through the analog crosspoint switch and configure an algorithm to perform the vibration analysis accordingly. As an example, sensors for capturing turbine cover vibration may be selected in the analog crosspoint switch to be passed on to a system that is configured with an algorithm to determine turbine cover vibration from the sensor signals. Upon completion of determining turbine cover vibration, the crosspoint switch may be configured to pass along thrust bearing axial vibration sensor signals and a corresponding vibration analysis algorithm may be applied to the data. In this way, each type of vibration may be analyzed by a single processing system that works cooperatively with an analog crosspoint switch to pass specific sensor signals for processing.

[0494] Referring to Figure 34, an analog crosspoint switch for collecting data in an industrial environment is depicted. The analog crosspoint switch 7022 may have a plurality of inputs 7024 that connect to sensors 7026 in the industrial environment. The analog crosspoint switch 7022 may also comprise a plurality of outputs 7028 that connect to data collection infrastructure, such as analog-to-digital converters 7030, analog comparators 7032, and the like. The analog crosspoint switch 7022 may facilitate connecting one or more inputs 7024 to one or more outputs 7028 by interpreting a switch control value that may be provided to it by a controller 7034 and the like.

[0495] An example system for data collection in an industrial environment comprising includes analog signal sources that each connect to at least one input of an analog crosspoint switch including a plurality of inputs and a plurality of outputs; where the analog crosspoint switch is configurable to switch a portion of the input signal sources to a plurality of the outputs.

[0496] 2. In certain example embodiments, the analog crosspoint switch further includes an analog-to-digital converter that converts a portion of analog signals input to the crosspoint switch

into representative digital signals; a portion of the outputs including analog outputs and a portion of the outputs comprises digital outputs; and/or where the analog crosspoint switch is adapted to detect one or more analog input signal conditions. Any one or more of the example embodiments include the analog input signal conditions including a voltage range of the signal, and where the analog crosspoint switch responsively adjusts input circuitry to comply with detected voltage range.

[0497] An example system of data collection in an industrial environment includes a number of industrial sensors that produce analog signals representative of a condition of an industrial machine in the environment being sensed by the number of industrial sensors, a crosspoint switch that receives the analog signals and routes the analog signals to separate analog outputs of the crosspoint switch based on a signal route plan presented to the crosspoint switch. In certain example embodiments, the analog crosspoint switch further includes an analog-to-digital converter that converts a portion of analog signals input to the crosspoint switch into representative digital signals; where a portion of the outputs include analog outputs and a portion of the outputs include digital outputs; where the analog crosspoint switch is adapted to detect one or more analog input signal conditions; where the one or more analog input signal conditions include a voltage range of the signal, and/or where the analog crosspoint switch responsively adjusts input circuitry to comply with detected voltage range.

[0498] An example method of data collection in an industrial environment includes routing a number of analog signals along a plurality of analog signal paths by connecting the plurality of analog signals individually to inputs of an analog crosspoint switch, configuring the analog crosspoint switch with data routing information from a data collection template for the industrial environment routing, and routing with the configured analog crosspoint switch a portion of the number of analog signals to a portion the plurality of analog signal paths. In certain further embodiments, at least one output of the analog crosspoint switch includes a high current driver circuit; at least one input of the analog crosspoint switch includes a voltage limiting circuit; and/or at least one input of the analog crosspoint switch includes a current limiting circuit. In certain further embodiments, the analog crosspoint switch includes a number of interconnected relays that facilitate connecting any of a number of inputs to any of a plurality of outputs; the analog crosspoint switch further including an analog-to-digital converter that converts a portion of analog signals input to the crosspoint switch into a representative digital signal; the analog crosspoint switch further including signal processing functionality to detect one or more analog input signal conditions, and in response thereto, to perform an action (e.g., set an alarm, change switch configuration, disable one or more outputs, power off a portion of the switch, change a state of a general purpose (digital / analog) output, etc.); where a portion of the outputs are analog outputs and a portion of the outputs are digital outputs; where the analog crosspoint switch is adapted to detect one or more analog input signal conditions; where the analog crosspoint switch is adapted to take one or more actions in response to detecting the one or more analog input signal conditions, the one more actions including setting an alarm, sending an alarm signal, changing a configuration of the analog crosspoint switch, disabling an output, powering off a portion of the analog crosspoint

switch, powering on a portion of the analog crosspoint switch, and/or controlling a general purpose output of the analog crosspoint switch.

[0499] An example system includes a power roller of a conveyor, including any of the described operations of an analog crosspoint switch. Without limitation, further example embodiments include sensing conditions of the power roller by the sensors to determine a rate of rotation of the power roller, a load being transported by the power roller, power being consumed by the power roller, and/or a rate of acceleration of the power roller. An example system includes a fan in a factory setting, including any of the described operations of an analog crosspoint switch. Without limitation, certain further embodiments include sensors disposed to sense conditions of the fan, including a fan blade tip speed, torque, back pressure, RPMs, and/or a volume of air per unit time displaced by the fan. An example system includes a turbine in a power generation environment, including any of the described operations of an analog crosspoint switch. Without limitation, certain further embodiments include a number of sensors disposed to sense conditions of the turbine, where the sensed conditions include a relative shaft vibration, an absolute vibration of bearings, a turbine cover vibration, a thrust bearing axial vibration, vibrations of stators or stator cores, vibrations of stator bars, and/or vibrations of stator end windings.

[0500] In example embodiments, methods and systems of data collection in an industrial environment may include a plurality of industrial condition sensing and acquisition modules that may include at least one programmable logic component per module that may control a portion of the sensing and acquisition functionality of its module. The programmable logic components on each of the modules may be interconnected by a dedicated logic bus that may include data and control channels. The dedicated logic bus may extend logically and/or physically to other programmable logic components on other sensing and acquisition modules. In example embodiments, the programmable logic components may be programmed via the dedicated interconnection bus, via a dedicated programming portion of the dedicated interconnection bus, via a program that is passed between programmable logic components, sensing and acquisition modules, or whole systems. A programmable logic component for use in an industrial environment data sensing and acquisition system may be a Complex Programmable Logic Device, an Application-Specific Integrated Circuit, microcontrollers, and combinations thereof.

[0501] A programmable logic component in an industrial data collection environment may perform control functions associated with data collection. Control examples include power control of analog channels, sensors, analog receivers, analog switches, portions of logic modules (e.g., a logic board, system, and the like) on which the programmable logic component is disposed, self-power-up/down, self-sleep/wake up, and the like. Control functions, such as these and others, may be performed in coordination with control and operational functions of other programmable logic components, such as other components on a single data collection module and components on other such modules. Other functions that a programmable logic component may provide may include generation of a voltage reference, such as a precise voltage reference for input signal condition detection. A programmable logic component may generate, set, reset, adjust, calibrate, or otherwise determine the voltage of the reference, its tolerance, and the like. Other functions of a

programmable logic component may include enabling a digital phase lock loop to facilitate tracking slowly transitioning input signals, and further to facilitate detecting the phase of such signals. Relative phase detection may also be implemented, including phase relative to trigger signals, other analog inputs, on-board references (e.g., on-board timers), and the like. A programmable logic component may be programmed to perform input signal peak voltage detection and control input signal circuitry, such as to implement auto-scaling of the input to an operating voltage range of the input. Other functions that may be programmed into a programmable logic component may include determining an appropriate sampling frequency for sampling inputs independently of their operating frequencies. A programmable logic component may be programmed to detect a maximum frequency among a plurality of input signals and set a sampling frequency for each of the input signals that is greater than the detected maximum frequency.

[0502] A programmable logic component may be programmed to configure and control data routing components, such as multiplexers, crosspoint switches, analog-to-digital converters, and the like, to implement a data collection template for the industrial environment. A data collection template may be included in a program for a programmable logic component. Alternatively, an algorithm that interprets a data collection template to configure and control data routing resources in the industrial environment may be included in the program.

[0503] In example embodiments, one or more programmable logic components in an industrial environment may be programmed to perform smart-band signal analysis and testing. Results of such analysis and testing may include triggering smart band data collection actions, that may include reconfiguring one or more data routing resources in the industrial environment. A programmable logic component may be configured to perform a portion of smart band analysis, such as collection and validation of signal activity from one or more sensors that may be local to the programmable logic component. Smart band signal analysis results from a plurality of programmable logic components may be further processed by other programmable logic components, servers, machine learning systems, and the like to determine compliance with a smart band.

[0504] In example embodiments, one or more programmable logic components in an industrial environment may be programmed to control data routing resources and sensors for outcomes, such as reducing power consumption (e.g., powering on/off resources as needed), implementing security in the industrial environment by managing user authentication, and the like. In example embodiments, certain data routing resources, such as multiplexers and the like, may be configured to support certain input signal types. A programmable logic component may configure the resources based on the type of signals to be routed to the resources. In example embodiments, the programmable logic component may facilitate coordination of sensor and data routing resource signal type matching by indicating to a configurable sensor a protocol or signal type to present to the routing resource. A programmable logic component may facilitate detecting a protocol of a signal being input to a data routing resource, such as an analog crosspoint switch and the like. Based on the detected protocol, the programmable logic component may configure routing resources to facilitate support and efficient processing of the protocol. In an example, a

programmable logic component configured data collection module in an industrial environment may implement an intelligent sensor interface specification, such as IEEE 1451.2 intelligent sensor interface specification.

[0505] Referring to Figure 35, an exemplary embodiment of a system for data collection in an industrial environment comprising distributed CPLDs interconnected by a bus for control and/or programming thereof is depicted. An exemplary data collection module 7200 may comprise one or more CPLDs 7206 for controlling one or more data collection system resources, such as sensors 7202 and the like. Other data collection resources that a CPLD may control may include crosspoint switches, multiplexers, data converters, and the like. CPLDs on a module may be interconnected by a bus, such as a dedicated logic bus 7204 that may extend beyond a data collection module to CPLDs on other data collection modules. Data collection modules, such as module 7200 may be configured in the environment, such as on an industrial machine 7208 (e.g., an impulse gas turbine) and/or 7210 (e.g., a co-generation system), and the like. Control and/or configuration of the CPLDs may be handled by a controller 7212 in the environment. Data collection and routing resources and interconnection (not shown) may also be configured within and among data collection modules 7200 as well as between and among industrial machines 7208 and 7210, and/or with external systems, such as Internet portals, data analysis servers, and the like to facilitate data collection, routing, storage, analysis, and the like.

[0506] An example system for data collection in an industrial environment includes a number of industrial condition sensing and acquisition modules, with a programmable logic component disposed on each of the modules, where the programmable logic component controls a portion of the sensing and acquisition functional of the corresponding module. The system includes communication bus that is dedicated to interconnecting the at least one programmable logic component disposed on at least one of the plurality of modules, wherein the communication bus extends to other programmable logic components on other sensing and acquisition modules.

[0507] In certain further embodiments, a system includes the programmable logic component programmed via the communication bus, the communication bus including a portion dedicated to programming of the programmable logic components, controlling a portion of the sensing and acquisition functionality of a module by a power control function such as: controlling power of a sensor, a multiplexer, a portion of the module, and/or controlling a sleep mode of the programmable logic component; controlling a portion of the sensing and acquisition functionality of a module by providing a voltage reference to a sensor and/or an analog-to-digital converter disposed on the module, by detecting a relative the phase of at least two analog signals derived from at least two sensors disposed on the module; by controlling sampling of data provided by at least one sensor disposed on the module; by detecting a peak voltage of a signal provided by a sensor disposed on the module; and/or by configuring at least one multiplexer disposed on the module by specifying to the multiplexer a mapping of at least one input and one output. In certain example embodiments, the communication bus extends to other programmable logic components on other condition sensing and/or acquisition modules. In certain example embodiments, a module may be an industrial environment condition sensing module. In certain example embodiments, a module

control program includes an algorithm for implementing an intelligent sensor interface communication protocol, such as an IEEE1451.2 compatible intelligent sensor interface communication protocol. In certain example embodiments, a programmable logic component includes configuring the programmable logic component and/or the sensing or acquisition module to implement a smart band data collection template. Example and non-limiting programmable logic components include field programmable gate arrays, complex programmable logic devices, and/or microcontrollers.

[0508] An example system includes a drilling machine for oil and gas field use, with a condition sensing and/or acquisition module to monitor aspects of a drilling machine. Without limitation, a further example system includes monitoring a compressor and/or monitoring an impulse steam engine.

[0509] In example embodiments, a system for data collection in an industrial environment may include a trigger signal and at least one data signal that share a common output of a signal multiplexer and upon detection of a condition in the industrial environment, such as a state of the trigger signal, the common output is switched to propagate either the data signal or the trigger signal. Sharing an output between a data signal and a trigger signal may also facilitate reducing a number of individually routed signals in an industrial environment. Benefits of reducing individually routed signals may include reducing the number of interconnections between data collection module, thereby reducing the complexity of the industrial environment. Trade-offs for reducing individually routed signals may include increasing sophistication of logic at signal switching modules to implement the detection and conditional switching of signals. A net benefit of this added localized logic complexity may be an overall reduction in the implementation complexity of such a data collection system in an industrial environment.

[0510] Exemplary deployment environments may include environments with trigger signal channel limitations, such as existing data collection systems that do not have separate trigger support for transporting an additional trigger signal to a module with sufficient computing sophistication to perform trigger detection. Another exemplary deployment may include systems that require at least some autonomous control for performing data collection.

[0511] In example embodiments, a system for data collection in an industrial environment may include an analog switch that switches between a first input, such as a trigger input and a second input, such as a data input based on a condition of the first input. A trigger input may be monitored by a portion of the analog switch to detect a change in the signal, such as from a lower voltage to a higher voltage relative to a reference or trigger threshold voltage. In example embodiments, a device that may receive the switched signal from the analog switch may monitor the trigger signal for a condition that indicates a condition for switching from the trigger input to the data input. When a condition of the trigger input is detected, the analog switch may be reconfigured, to direct the data input to the same output that was propagating the trigger output.

[0512] In example embodiments, a system for data collection in an industrial environment may include an analog switch that directs a first input to an output of the analog switch until such time as the output of the analog switch indicates that a second input should be directed to the output of

the analog switch. The output of the analog switch may propagate a trigger signal to the output. In response to the trigger signal propagating through the switch transitioning from a first condition (e.g., a first voltage below a trigger threshold voltage value) to a second condition (e.g., a second voltage above the trigger threshold voltage value), the switch may stop propagating the trigger signal and instead propagate another input signal to the output. In example embodiments, the trigger signal and the other data signal may be related, such as the trigger signal may indicate a presence of an object being placed on a conveyer and the data signal represents a strain placed on the conveyer.

[0513] In example embodiments, to facilitate timely detection of the trigger condition, a rate of sampling of the output of the analog switch may be adjustable, so that, for example, the rate of sampling is higher while the trigger signal is propagated and lower when the data signal is propagated. Alternatively, a rate of sampling may be fixed for either the trigger or the data signal. In example embodiments, the rate of sampling may be based on a predefined time from trigger occurrence to trigger detection and may be faster than a minimum sample rate to capture the data signal.

[0514] In example embodiments, routing a plurality of hierarchically organized triggers onto another analog channel may facilitate implementing a hierarchical data collection triggering structure in an industrial environment. A data collection template to implement a hierarchical trigger signal architecture may include signal switch configuration and function data that may facilitate a signal switch facility, such as an analog crosspoint switch or multiplexer to output a first input trigger in a hierarchy, and based on the first trigger condition being detected, output a second input trigger in the hierarchy on the same output as the first input trigger by changing an internal mapping of inputs to outputs. Upon detection of the second input trigger condition, the output may be switched to a data signal, such as data from a sensor in an industrial environment.

[0515] In example embodiments, upon detection of a trigger condition, in addition to switching from the trigger signal to a data signal, an alarm may be generated and optionally propagated to a higher functioning device/module. In addition to switching to a data signal, upon detection of a state of the trigger, sensors that otherwise may be disabled or powered down may be energized/activated to begin to produce data for the newly selected data signal. Activating might alternatively include sending a reset or refresh signal to the sensor(s).

[0516] In example embodiments, a system for data collection in an industrial environment may include a system for routing a trigger signal onto a data signal path in association with a gearbox of an industrial vehicle. Combining a trigger signal onto a signal path that is also used for a data signal may be useful in gearbox applications by reducing the number of signal lines that need to be routed, while enabling advanced functions, such as data collection based on pressure changes in the hydraulic fluid and the like. As an example, a sensor may be configured to detect a pressure difference in the hydraulic fluid that exceeds a certain threshold as may occur when the hydraulic fluid flow is directed back into the impeller to give higher torque at low speeds. The output of such a sensor may be configured as a trigger for collecting data about the gearbox when operating at low speeds. In an example, a data collection system for an industrial environment may have a

multiplexer or switch that facilitates routing either a trigger or a data channel over a single signal path. Detecting the trigger signal from the pressure sensor may result in a different signal being routed through the same line that the trigger signal was routed by switching a set of controls. A multiplexer may, for example, output the trigger signal until the trigger signal is detected as indicating that the output should be changed to the data signal. As a result of detecting the high-pressure condition, a data collection activity may be activated so that data can be collected using the same line that was recently used by the trigger signal.

[0517] In example embodiments, a system for data collection in an industrial environment may include a system for routing a trigger signal onto a data signal path in association with a vehicle suspension for truck and car operation. Vehicle suspension, particularly active suspension may include sensors for detecting road events, suspension conditions, and vehicle data, such as speed, steering, and the like. These conditions may not always need to be detected, except, for example, upon detection of a trigger condition. Therefore, combining the trigger condition signal and at least one data signal on a single physical signal routing path could be implemented. Doing so may reduce costs due to fewer physical connections required in such a data collection system. In an example, a sensor may be configured to detect a condition, such as a pothole, to which the suspension must react. Data from the suspension may be routed along the same signal routing path as this road condition trigger signal so that upon detection of the pothole, data may be collected that may facilitate determining aspects of the suspension's reaction to the pothole.

[0518] In example embodiments, a system for data collection in an industrial environment may include a system for routing a trigger signal onto a data signal path in association with a turbine for power generation in a power station. A turbine used for power generation may be retrofitted with a data collection system that optimizes existing data signal lines to implement greater data collection functions. One such approach involves routing new sources of data over existing lines. While multiplexing signals generally satisfies this need, combining a trigger signal with a data signal via a multiplexer or the like can further improve data collection. In an example, a first sensor may include a thermal threshold sensor that may measure the temperature of an aspect of a power generation turbine. Upon detection of that trigger (e.g., by the temperature rising above the thermal threshold), a data collection system controller may send a different data collection signal over the same line that was used to detect the trigger condition. This may be accomplished by a controller or the like sensing the trigger signal change condition and then signaling to the multiplexer to switch from the trigger signal to a data signal to be output on the same line as the trigger signal for data collection. In this example, when a turbine is detected as having a portion that exceeds its safe thermal threshold, a secondary safety signal may be routed over the trigger signal path and monitored for additional safety conditions, such as overheating and the like.

[0519] Referring to Figure 36, an embodiment of routing a trigger signal over a data signal path in a data collection system in an industrial environment is depicted. Signal multiplexer 7400 may receive a trigger signal on a first input from a sensor or other trigger source 7404 and a data signal on a second input from a sensor for detecting a temperature associated with an industrial machine in the environment 7402. The multiplexer 7400 may be configured to output the trigger signal onto

an output signal path 7406. A data collection module 7410 may process the signal on the data path 7406 looking for a change in the signal indicative of a trigger condition provided from the trigger sensor 7404 through the multiplexer 7400. Upon detection, a control output 7408 may be changed and thereby control the multiplexer 7400 to start outputting data from the temperature probe 7402 by switching an internal switch or the like that may control one or more of the inputs that may be routed to the output 7406. Data collection facility 7410 may activate a data collection template in response to the detected trigger that may include switching the multiplexer and collecting data into triggered data storage 7412. Upon completion of the data collection activity, multiplexer control signal 7408 may revert to its initial condition so that trigger sensor 7404 may be monitored again.

[0520] An example system for data collection in an industrial environment includes an analog switch that directs a first input to an output of the analog switch until such time as the output of the analog switch indicates that a second input should be directed to the output of the analog switch. In certain further embodiments, the example system includes: where the output of the analog switch indicated that the second input should be directed to the output based on the output transitioning from a pending condition to a triggered condition; wherein the triggered condition includes detecting the output presenting a voltage above a trigger voltage value; routing a number of signals with the analog switch from inputs on the analog switch to outputs on the analog switch in response to the output of the analog switch indicating that the second input should be directed to the output; sampling the output of the analog switch at a rate that exceeds a rate of transition for a number of signals input to the analog switch; and/or generating an alarm signal when the output of the analog switch indicates that a second input should be directed to the output of the analog switch.

[0521] An example system for data collection in an industrial environment includes an analog switch that switches between a first input and a second input based on a condition of the first input. In certain further embodiments, the condition of the first input comprises the first input presenting a triggered condition, and/or the triggered condition includes detecting the first input presenting a voltage above a trigger voltage value. In certain example embodiments, the analog switch includes routing a plurality of signals with the analog from inputs on the analog switch to outputs on the analog switch based on the condition of the first input, sampling an input of the analog switch at a rate that exceeds a rate of transition for a plurality of signals input to the analog switch, and/or generating an alarm signal based on the condition of the first input.

[0522] An example system for data collection in an industrial environment includes a trigger signal and at least one data signal that share a common output of a signal multiplexer, and upon detection of a predefined state of the trigger signal, the common output is configured to propagate the at least one data signal through the signal multiplexer. In certain further embodiments, the signal multiplexer is an analog multiplexer, the predefined state of the trigger signal is detected on the common output, detection of the predefined state of the trigger signal includes detecting the common output presenting a voltage above a trigger voltage value, the multiplexer includes routing a plurality of signals with the multiplexer from inputs on the multiplexer to outputs on the multiplexer in response to detection of the predefined state of the trigger signal, the multiplexer includes sampling the output of the multiplexer at a rate that exceeds a rate of transition for a

plurality of signals input to the multiplexer, the multiplexer includes generating an alarm in response to detection of the predefined state of the trigger signal, and/or the multiplexer includes activating at least one sensor to produce the at least one data signal. Without limitation, example systems include: monitoring a gearbox of an industrial vehicle by directing a trigger signal representing a condition of the gearbox to an output of the analog switch until such time as the output of the analog switch indicates that a second input representing a condition of the gearbox related to the trigger signal should be directed to the output of the analog switch; monitoring a suspension system of an industrial vehicle by directing a trigger signal representing a condition of the suspension to an output of the analog switch until such time as the output of the analog switch indicates that a second input representing a condition of the suspension related to the trigger signal should be directed to the output of the analog switch; and/or monitoring a power generation turbine by directing a trigger signal representing a condition of the power generation turbine to an output of the analog switch until such time as the output of the analog switch indicates that a second input representing a condition of the power generation turbine related to the trigger signal should be directed to the output of the analog switch.

[0523] In example embodiments, a system for data collection in an industrial environment may include a data collection system that monitors at least one signal for a set of collection band parameters and upon detection of a parameter from the set of collection band parameters in the signal, configures collection of data from a set of sensors based on the detected parameter. The set of selected sensors, the signal, and the set of collection band parameters may be part of a smart bands data collection template that may be used by the system when collecting data in an industrial environment. A motivation for preparing a smart-bands data collection template may include monitoring a set of conditions of an industrial machine to facilitate improved operation, reduce down time, preventive maintenance, failure prevention, and the like. Based on analysis of data about the industrial machine, such as those conditions that may be detected by the set of sensors, an action may be taken, such as notifying a user of a change in the condition, adjusting operating parameters, scheduling preventive maintenance, triggering data collection from additional sets of sensors, and the like. An example of data that may indicate a need for some action may include changes that may be detectable through trends present in the data from the set of sensors. Another example is trends of analysis values derived from the set of sensors.

[0524] In example embodiments, the set of collection band parameters may include values received from a sensor that is configured to sense a condition of the industrial machine (e.g., bearing vibration). However, a set of collection band parameters may instead be a trend of data received from the sensor (e.g., a trend of bearing vibration across a plurality of vibration measurements by a bearing vibration sensor). In example embodiments, a set of collection band parameters may be a composite of data and/or trends of data from a plurality of sensors (e.g., a trend of data from on-axis and off-axis vibration sensors). In example embodiments, when a data value derived from one or more sensors as described herein is sufficiently close to a value of data in the set of collection band parameters, the data collection activity from the set of sensors may be triggered. Alternatively, a data collection activity from the set of sensors may be triggered when a

data value derived from the one or more sensors (e.g., trends and the like) falls outside of a set of collection band parameters. In an example, a set of data collection band parameters for a motor may be a range of rotational speeds from 95% to 105% of a select operational rotational speed. So long as a trend of rotational speed of the motor stays within this range, a data collection activity may be deferred. However, when the trend reaches or exceeds this range, then a data collection activity, such as one defined by a smart bands data collection template may be triggered.

[0525] In example embodiments, triggering a data collection activity, such as one defined by a smart bands data collection template, may result in a change to a data collection system for an industrial environment that may impact aspects of the system such as data sensing, switching, routing, storage allocation, storage configuration, and the like. This change to the data collection system may occur in near real time to the detection of the condition; however, it may be scheduled to occur in the future. It may also be coordinated with other data collection activities so that active data collection activities, such as a data collection activity for a different smart bands data collection template, can complete prior to the system being reconfigured to meet the smart bands data collection template that is triggered by the sensed condition meeting the smart bands data collection trigger.

[0526] In example embodiments, processing of data from sensors may be cumulative over time, over a set of sensors, across machines in an industrial environment, and the like. While a sensed value of a condition may be sufficient to trigger a smart bands data collection template activity, data may need to be collected and processed over time from a plurality of sensors to generate a data value that may be compared to a set of data collection band parameters for conditionally triggering the data collection activity. Using data from multiple sensors and/or processing data, such as to generate a trend of data values and the like may facilitate preventing inconsequential instances of a sensed data value being outside of an acceptable range from causing unwarranted smart bands data collection activity. In an example, if a vibration from a bearing is detected outside of an acceptable range infrequently, then trending for this value over time may be useful to detect if the frequency is increasing, decreasing, or staying substantially constant or within a range of values. If the frequency of such a value is found to be increasing, then such a trend is indicative of changes occurring in the operation of the industrial machine as experienced by the bearing. An acceptable range of values of this trended vibration value may be established as a set of data collection band parameters against which vibration data for the bearing will be monitored. When the trended vibration value is outside of this range of acceptable values, a smart bands data collection activity may be activated.

[0527] In example embodiments, a system for data collection in an industrial environment that supports smart band data collection templates may be configured with data processing capability at a point of sensing of one or more conditions that may trigger a smart bands data collection template data collection activity, such as: by use of an intelligent sensor that may include data processing capabilities; by use of a programmable logic component that interfaces with a sensor and processes data from the sensor; by use of a computer processor, such as a microprocessor and the like disposed proximal to the sensor; and the like. In example embodiments, processing of data

collected from one or more sensors for detecting a smart bands template data collection activity may be performed by remote processors, servers, and the like that may have access to data from a plurality of sensors, sensor modules, industrial machines, industrial environments, and the like.

[0528] In example embodiments, a system for data collection in an industrial environment may include a data collection system that monitors an industrial environment for a set of parameters, and upon detection of at least one parameter, configures the collection of data from a set of sensors and causes a data storage controller to adapt a configuration of data storage facilities to support collection of data from the set of sensors based on the detected parameter. The methods and systems described herein for conditionally changing a configuration of a data collection system in an industrial environment to implement a smart bands data collection template may further include changes to data storage architectures. As an example, a data storage facility may be disposed on a data collection module that may include one or more sensors for monitoring conditions in an industrial environment. This local data storage facility may typically be configured for rapid movement of sensed data from the module to a next level sensing or processing module or server. When a smart bands data collection condition is detected, sensor data from a plurality of sensors may need to be captured concurrently. To accommodate this concurrent collection, the local memory may be reconfigured to capture data from each of the plurality of sensors in a coordinated manner, such as repeatedly sampling each of the sensors synchronously, or with a known offset, and the like, to build up a set of sensed data that may be much larger than would typically be captured and moved through the local memory. A storage control facility for controlling the local storage may monitor the movement of sensor data into and out of the local data storage, thereby ensuring safe movement of data from the plurality of sensors to the local data storage and on to a destination, such as a server, networked storage facility, and the like. The local data storage facility may be configured so that data from the set of sensors associated with a smart bands data collection template are securely stored and readily accessible as a set of smart band data to facilitate processing the smart band-specific data. As an example, local storage may comprise non-volatile memory (NVM). To prepare for data collection in response to a smart band data collection template being triggered, portions of the NVM may be erased to prepare the NVM to receive data as indicated in the template.

[0529] In example embodiments, multiple sensors may be arranged into a set of sensors for condition-specific monitoring. Each set, which may be a logical set of sensors, may be selected to provide information about elements in an industrial environment that may provide insight into potential problems, root causes of problems, and the like. Each set may be associated with a condition that may be monitored for compliance with an acceptable range of values. The set of sensors may be based on a machine architecture, hierarchy of components, or a hierarchy of data that contributes to a finding about a machine that may usefully be applied to maintaining or improving performance in the industrial environment. Smart band sensor sets may be configured based on expert system analysis of complex conditions, such as machine failures and the like. Smart band sensor sets may be arranged to facilitate knowledge gathering independent of a particular failure mode or history. Smart band sensor sets may be arranged to test a suggested smart

band data collection template prior to implementing it as part of an industrial machine operations program. Gathering and processing data from sets of sensors may facilitate determining which sensors contribute meaningful data to the set, and those sensors that do not contribute can be removed from the set. Smart band sensor sets may be adjusted based on external data, such as industry studies that indicate the types of sensor data that are most helpful to reduce failures in an industrial environment.

[0530] In example embodiments, a system for data collection in an industrial environment may include a data collection system that monitors at least one signal for compliance to a set of collection band conditions and upon detection of a lack of compliance, configures the collection of data from a predetermined set of sensors associated with the monitored signal. Upon detection of a lack of compliance, a collection band template associated with the monitored signal may be accessed, and resources identified in the template may be configured to perform the data collection. In example embodiments, the template may identify sensors to activate, data from the sensors to collect, duration of collection or quantity of data to be collected, destination (e.g., memory structure) to store the collected data, and the like. In example embodiments, a smart band method for data collection in an industrial environment may include periodic collection of data from one or more sensors configured to sense a condition of an industrial machine in the environment. The collected data may be checked against a set of criteria that define an acceptable range of the condition. Upon validation that the collected data is either approaching one end of the acceptable limit or is beyond the acceptable range of the condition, data collection may commence from a smart-band group of sensors associated with the sensed condition based on a smart-band collection protocol configured as a data collection template. In example embodiments, an acceptable range of the condition is based on a history of applied analytics of the condition. In example embodiments, upon validation of the acceptable range being exceeded, data storage resources of a module in which the sensed condition is detected may be configured to facilitate capturing data from the smart band group of sensors.

[0531] In example embodiments, monitoring a condition to trigger a smart band data collection template data collection action may be: in response to: a regulation, such as a safety regulation; in response to an upcoming activity, such as a portion of the industrial environment being shut down for preventive maintenance; in response to sensor data missing from routine data collection activities; and the like. In example embodiments, in response to a faulty sensor or sensor data missing from a smart band template data collection activity, one or more alternate sensors may be temporarily included in the set of sensors so as to provide data that may effectively substitute for the missing data in data processing algorithms.

[0532] In example embodiments, smart band data collection templates may be configured for detecting and gathering data for smart band analysis covering vibration spectra, such as vibration envelope and current signature for spectral regions or peaks that may be combinations of absolute frequency or factors of machine related parameters, vibration time waveforms for time-domain derived calculations including, without limitation: RMS overall, peak overall, true peak, crest factor, and the like; vibration vectors, spectral energy humps in various regions (e.g., low-

frequency region, high frequency region, low orders, and the like); pressure-volume analysis and the like.

[0533] In example embodiments, a system for data collection that applies smart band data collection templates may be applied to an industrial environment, such as ball screw actuators in an automated production environment. Smart band analysis may be applied to ball screw actuators in industrial environments such as precision manufacturing or positioning applications (e.g., semiconductor photolithography machines, and the like). As a typical primary objective of using a ball screw is for precise positioning, detection of variation in the positioning mechanism can help avoid costly defective production runs. Smart bands triggering and data collection may help in such applications by detecting, through smart band analysis, potential variations in the positioning mechanism such as in the ball screw mechanism, a worm drive, a linear motor, and the like. In an example, data related to a ball screw positioning system may be collected with a system for data collection in an industrial environment as described herein. A plurality of sensors may be configured to collect data such as screw torque, screw direction, screw speed, screw step, screw home detection, and the like. Some portion of this data may be processed by a smart bands data analysis facility to determine if variances, such as trends in screw speed as a function of torque, approach or exceed an acceptable threshold. Upon such a determination, a data collection template for the ball screw production system may be activated to configure the data sensing, routing, and collection resources of the data collection system to perform data collection to facilitate further analysis. The smart band data collection template facilitates rapid collection of data from other sensors than screw speed and torque, such as position, direction, acceleration, and the like by routing data from corresponding sensors over one or more signal paths to a data collector. The duration and order of collection of the data from these sources may be specified in the smart bands data collection template so that data required for further analysis is effectively captured.

[0534] In example embodiments, a system for data collection that applies smart band data collection templates to configure and utilize data collection and routing infrastructure may be applied to ventilation systems in mining environments. Ventilation provides a crucial role in mining safety. Early detection of potential problems with ventilation equipment can be aided by applying a smart bands approach to data collection in such an environment. Sensors may be disposed for collecting information about ventilation operation, quality, and performance throughout a mining operation. At each ventilation device, ventilation-related elements, such as fans, motors, belts, filters, temperature gauges, voltage, current, air quality, poison detection, and the like may be configured with a corresponding sensor. While variation in any one element (e.g., air volume per minute, and the like) may not be indicative of a problem, smart band analysis may be applied to detect trends over time that may be suggestive of potential problems with ventilation equipment. To perform smart bands analysis, data from a plurality of sensors may be required to form a basis for analysis. By implementing data collection systems for ventilation stations, data from a ventilation system may be captured. In an example, a smart band analysis may be indicated for a ventilation station. In response to this indication, a data collection system may be configured

to collect data by routing data from sensors disposed at the ventilation station to a central monitoring facility that may gather and analyze data from several ventilation stations.

[0535] In example embodiments, a system for data collection that applies smart band data collection templates to configure and utilize data collection and routing infrastructure may be applied to drivetrain data collection and analysis in mining environments. A drivetrain, such as a drivetrain for a mining vehicle, may include a range of elements that could benefit from use of the methods and systems of data collection in an industrial environment as described herein. In particular, smart band-based data collection may be used to collect data from heavy duty mining vehicle drivetrains under certain conditions that may be detectable by smart bands analysis. A smart bands-based data collection template may be used by a drivetrain data collection and routing system to configure sensors, data paths, and data collection resources to perform data collection under certain circumstances, such as those that may indicate an unacceptable trend of drivetrain performance. A data collection system for an industrial drivetrain may include sensing aspects of a non-steering axle, a planetary steering axle, driveshafts, (e.g., main and wing shafts), transmissions, (e.g., standard, torque converters, long drop), and the like. A range of data related to these operational parts may be collected. However, data for support and structural members that support the drivetrain may also need to be collected for thorough smart band analysis. Therefore, collection across this wide range of drivetrain-related components may be triggered based on a smart band analysis determination of a need for this data. In an example, a smart band analysis may indicate potential slippage between a main and wing driveshaft that may be represented by an increasing trend in response delay time of the wing drive shaft to main drive shaft operation. In response to this increasing trend, data collection modules disposed throughout the mining vehicle's drive train may be configured to route data from local sensors to be collected and analyzed by data collectors. Mining vehicle drivetrain smart based data collection may include a range of templates based on which type of trend is detected. If a trend related to a steering axle is detected, a data collection template to be implemented may be different in sensor content, duration, and the like than for a trend related to power demand for a normalized payload. Each template could configure data sensing, routing, and collection resources throughout the vehicle drive train accordingly.

[0536] Referring to Figure 37, a system for data collection in an industrial environment that facilitates data collection for smart band analysis is depicted. A system for data collection in an industrial environment may include a smart band analysis data collection template repository 7600 in which smart band templates 7610 for data collection system configuration and collection of data may be stored and accessed by a data collection controller 7602. The templates 7610 may include data collection system configuration 7604 and operation information 7606 that may identify sensors, collectors, signal paths, and information for initiation and coordination of collection, and the like. A controller 7602 may receive an indication, such as a command from a smart band analysis facility 7608 to select and implement a specific smart band template 7610. The controller 7602 may access the template 7610 and configure the data collection system resources based on the information in that template. In example embodiments, the template may identify: specific sensors; a multiplexer/switch configuration, data collection trigger/initiation signals and/or

conditions, time duration and/or amount of data for collection; destination of collected data; intermediate processing, if any; and any other useful information, (e.g., instance identifier, and the like). The controller 7602 may configure and operate the data collection system to perform the collection for the smart band template and optionally return the system configuration to a previous configuration.

[0537] An example system for data collection in an industrial environment includes a data collection system that monitors at least one signal for a set of collection band parameters and, upon detection of a parameter from the set of collection band parameters, configures portions of the system and performs collection of data from a set of sensors based on the detected parameter. In certain further embodiments, the signal includes an output of a sensor that senses a condition in the industrial environment, where the set of collection band parameters comprises values derivable from the signal that are beyond an acceptable range of values derivable from the signal; where the at least one signal includes an output of a sensor that senses a condition in the industrial environment; wherein configuring portions of the system includes configuring a storage facility to accept data collected from the set of sensors; where configuring portions of the system includes configuring a data routing portion includes at least one of: an analog crosspoint switch, a hierarchical multiplexer, an analog-to-digital converter, an intelligent sensor, and/or a programmable logic component; wherein detection of a parameter from the set of collection band parameters comprises detecting a trend value for the signal being beyond an acceptable range of trend values; and/or where configuring portions of the system includes implementing a smart band data collection template associated with the detected parameter. In certain example embodiments, a data collection system monitors a signal for data values within a set of acceptable data values that represent acceptable collection band conditions for the signal and, upon detection of a data value for the at least one signal outside of the set of acceptable data values, triggers a data collection activity that causes collecting data from a predetermined set of sensors associated with the monitored signal. In certain further embodiment, a data collection system includes the signal including an output of a sensor that senses a condition in the industrial environment; where the set of acceptable data value includes values derivable from the signal that are within an acceptable range of values derivable from the signal; configuring a storage facility of the system to facilitate collecting data from the predetermined set of sensors in response to the detection of a data value outside of the set of acceptable data values; configuring a data routing portion of the system including an analog crosspoint switch, a hierarchical multiplexer, an analog-to-digital converter, an intelligent sensor, and/or a programmable logic component in response to detecting a data value outside of the set of acceptable data values; where detection of a data value for the signal outside of the set of acceptable data values includes detecting a trend value for the signal being beyond an acceptable range of trend values; and/or where the data collection activity is defined by a smart band data collection template associated with the detected parameter.

[0538] An example method for data collection in an industrial environment comprising includes an operation to collect data from sensor(s) configured to sense a condition of an industrial machine in the environment; an operation to check the collected data against a set of criteria that define an

acceptable range of the condition; and in response to the collected data violating the acceptable range of the condition, an operation to collect data from a smart-band group of sensors associated with the sensed condition based on a smart-band collection protocol configured as a smart band data collection template. In certain further embodiments, a method includes where violating the acceptable range of the condition includes a trend of the data from the sensor(s) approaching a maximum value of the acceptable range; where the smart-band group of sensors is defined by the smart band data collection template; where the smart band data collection template includes a list of sensors to activate, data from the sensors to collect, duration of collection of data from the sensors, and/or a destination location for storing the collected data; where collecting data from a smart-band group of sensors includes configuring at least one data routing resource of the industrial environment that facilitates routing data from the smart band group of sensors to a plurality of data collectors; and/or where the set of criteria includes a range of trend values derived by processing the data from sensor(s).

[0539] Without limitation, an example system monitors a ball screw actuator in an automated production environment, and monitors at least one signal from the ball screw actuator for a set of collection band parameters and, upon detection of a parameter from the set of collection band parameters, configures portions of the system and performs collection of data from a set of sensors disposed to monitor conditions of the ball screw actuator based on the detected parameter; another example system monitors a ventilation system in a mining environment, and monitors at least one signal from the ventilation system for a set of collection band parameters and, upon detection of a parameter from the set of collection band parameters, configures portions of the system and performs collection of data from a set of sensors disposed to monitor conditions of the ventilation system based on the detected parameter; an example system monitors a drivetrain of a mining vehicle, and monitors at least one signal from the drive train for a set of collection band parameters and, upon detection of a parameter from the set of collection band parameters, configures portions of the system and performs collection of data from a set of sensors disposed to monitor conditions of the drivetrain based on the detected parameter.

[0540] In example embodiments, a system for data collection in an industrial environment may automatically configure local and remote data collection resources and may perform data collection from a plurality of system sensors that are identified as part of a group of sensors that produce data that is required to perform operational deflection shape rendering. In example embodiments, the system sensors are distributed throughout structural portions of an industrial machine in the industrial environment. In example embodiments, the system sensors sense a range of system conditions including vibration, rotation, balance, friction, and the like. In example embodiments, automatically configuring is in response to a condition in the environment being detected outside of an acceptable range of condition values. In example embodiments, a sensor in the identified group of system sensors senses the condition.

[0541] In example embodiments, a system for data collection in an industrial environment may configure a data collection plan, such as a template, to collect data from a plurality of system sensors distributed throughout a machine to facilitate automatically producing an operational

deflection shape visualization ("ODSV") based on machine structural information and a data set used to produce an ODSV of the machine.

[0542] In example embodiments, a system for data collection in an industrial environment may configure a data collection template for collecting data in an industrial environment by identifying sensors disposed for sensing conditions of preselected structural members of an industrial machine in the environment based on an ODSV of the industrial machine. In example embodiments, the template may include an order and timing of data collection from the identified sensors.

[0543] In example embodiments, methods and systems for data collection in an industrial environment may include a method of establishing an acceptable range of sensor values for a plurality of industrial machine condition sensors by validating an operational deflection shape visualization of structural elements of the machine as exhibiting deflection within an acceptable range, wherein data from the plurality of sensors used in the validated ODSV define the acceptable range of sensor values.

[0544] In example embodiments, a system for data collection in an industrial environment may include a plurality of data sources, such as sensors, that may be grouped for coordinated data collection to provide data required to produce an ODSV. Information regarding the sensors to group, data collection coordination requirements, and the like may be retrieved from an ODSV data collection template. Coordinated data collection may include concurrent data collection. To facilitate concurrent data collection from a portion of the group of sensors, sensor routing resources of the system for data collection may be configured, such as by configuring a data multiplexer to route data from the portion of the group of sensors to which it connects to data collectors. In example embodiments, each such source that connects an input of the multiplexer may be routed within the multiplexer to separate outputs so that data from all of the connected sources may be routed on to data collection elements of the industrial environment. In example embodiments, the multiplexer may include data storage capabilities that may facilitate sharing a common output for at least a portion of the inputs. In example embodiments, a multiplexer may include data storage capabilities and data bus-enabled outputs so that data for each source may be captured in a memory and transmitted over a data bus, such as a data bus that is common to the outputs of the multiplexer. In example embodiments, sensors may be smart sensors that may include data storage capabilities and may send data from the data storage to the multiplexer in a coordinated manner that supports use of a common output of the multiplexer and/or use of a common data bus.

[0545] In example embodiments, a system for data collection in an industrial environment may comprise templates for configuring the data collection system to collect data from a plurality of sensors to perform ODSV for a plurality of deflection shapes. Individual templates may be configured for visualization of looseness, soft joints, bending, twisting, and the like. Individual deflection shape data collection templates may be configured for different portions of a machine in an industrial environment.

[0546] In example embodiments, a system for data collection in an industrial environment may facilitate operational deflection shape visualization that may include visualization of locations of sensors that contributed data to the visualization. In the visualization, each sensor that contributed

data to generate the visualization may be indicated by a visual element. The visual element may facilitate user access to information about the sensor, such as location, type, representative data contributed, path of data from the sensor to a data collector, a deflection shape template identifier, a configuration of a switch or multiplexer through which the data is routed, and the like. The visual element may be determined by associating sensor identification information received from a sensor with information, such as a sensor map, that correlates sensor identification information with physical location in the environment. The information may appear in the visualization in response to the visual element representing the sensor being selected, such as by a user positioning a cursor on the sensor visual element.

[0547] In example embodiments, ODSV may benefit from data satisfying a phase relationship requirement. A data collection system in the environment may be configured to facilitate collecting data that complies with the phase relationship requirement. Alternatively, the data collection system may be configured to collect data from a plurality of sensors that contains data that satisfies the phase relationship requirements but may also include data that does not. A post processing operation that may access phase detection data may select a subset of the collected data.

[0548] In example embodiments, a system for data collection in an industrial environment may include a multiplexer receiving data from a plurality of sensors and multiplexing the received data for delivery to a data collector. The data collector may process the data to facilitate ODSV. ODSV may require data from several different sensors, and may benefit from using a reference signal, such as data from a sensor, when processing data from the different sensors. The multiplexer may be configured to provide data from the different sensors, such as by switching among its inputs over time so that data from each sensor may be received by the data collector. However, the multiplexer may include a plurality of outputs so that at least a portion of the inputs may be routed to least two of the plurality of outputs. Therefore, in example embodiments, a multiple output multiplexer may be configured to facilitate data collection that may be suitable for ODSV by routing a reference signal from one of its inputs (e.g., data from an accelerometer) to one of its outputs and multiplexing data from a plurality of its outputs onto one or more of its outputs while maintaining the reference signal output routing. A data collector may collect the data from the reference output and use that to align the multiplexed data from the other sensors.

[0549] In example embodiments, a system for data collection in an industrial environment may facilitate ODSV through coordinated data collection related to conveyors for mining applications. Mining operations may rely on conveyor systems to move material, supplies, and equipment into and out of a mine. Mining operations may typically operate around the clock; therefore, conveyor downtime may have a substantive impact on productivity and costs. Advanced analysis of conveyor and related systems that focuses on secondary affects that may be challenging to detect merely through point observation may be more readily detected via ODSV. Capturing operational data related to vibration, stresses, and the like can facilitate ODSV. However, coordination of data capture provides more reliable results. Therefore, a data collection system that may have sensors dispersed throughout a conveyor system can be configured to facilitate such coordinated data collection. In an example, capture of data affecting structural components of a conveyor, such as;

landing points and the horizontal members that connect them and support the conveyer between landing points; conveyer segment handoff points; motor mounts; mounts of conveyer rollers and the like may need to be coordinated with data related to conveyor dynamic loading, drive systems, motors, gates, and the like. A system for data collection in an industrial environment, such as a mining environment may include data sensing and collection modules placed throughout the conveyor at locations such as segment handoff points, drive systems, and the like. Each module may be configured by one or more controllers, such as programmable logic controllers, that may be connected through a physical or logical (e.g., wireless) communication bus that aids in performing coordinated data collection. To facilitate coordination, a reference signal, such as a trigger and the like, may be communicated among the modules for use when collecting data. In example embodiments, data collection and storage may be performed at each module so as to reduce the need for real-time transfer of sensed data throughout the mining environment. Transfer of data from the modules to an ODSV processing facility may be performed after collection, or as communication bandwidth between the modules and the processing facility allows. ODSV can provide insight into conditions in the conveyer, such as deflection of structural members that may, over time cause premature failure. Coordinated data collection with a data collection system for use in an industrial environment, such as mining, can enable ODSV that may reduce operating costs by reducing downtime due to unexpected component failure.

[0550] In example embodiments, a system for data collection in an industrial environment may facilitate operational deflection shape visualization through coordinated data collection related to fans for mining applications. Fans provide a crucial function in mining operations of moving air throughout a mine to provide ventilation, equipment cooling, combustion exhaust evacuation, and the like. Ensuring reliable and often continuous operation of fans may be critical for miner safety and cost-effective operations. Dozens or hundreds of fans may be used in large mining operations. Fans, such as fans for ventilation management may include circuit, booster, and auxiliary types. High capacity auxiliary fans may operate at high speeds, over 2500 RPMs. Performing ODSV may reveal important reliability information about fans deployed in a mining environment. Collecting the range of data needed for ODSV of mining fans may be performed by a system for collecting data in industrial environments as described herein. In example embodiments, sensing elements, such as intelligent sensing and data collection modules may be deployed with fans and/or fan subsystems. These modules may exchange collection control information (e.g., over a dedicated control bus and the like) so that data collection may be coordinated in time and phase to facilitate ODSV.

[0551] A large auxiliary fan for use in mining may be constructed for transportability into and through the mine and therefore may include a fan body, intake and outlet ports, dilution valves, protection cage, electrical enclosure, wheels, access panels, and other structural and/or operational elements. The ODSV of such an auxiliary fan may require collection of data from many different elements. A system for data collection may be configured to sense and collect data that may be combined with structural engineering data to facilitate ODSV for this type of industrial fan.

[0552] Referring to Figure 38, an embodiment of a system for data collection in an industrial environment that performs coordinated data collection suitable for ODSV is depicted. A system for data collection in an industrial environment may include a ODSV data collection template repository 7800 in which ODSV templates 7810 for data collection system configuration and collection of data may be stored and accessed by a system for data collection controller 7802. The templates 7810 may include data collection system configuration 7804 and operation information 7806 that may identify sensors, collectors, signal paths, reference signal information, information for initiation and coordination of collection, and the like. A controller 7802 may receive an indication, such as a command from an ODSV analysis facility 7808 to select and implement a specific ODSV template 7810. The controller 7802 may access the template 7810 and configure the data collection system resources based on the information in that template. In example embodiments, the template may identify specific sensors, multiplexer/switch configuration, reference signals for coordinating data collection, data collection trigger/initiation signals and/or conditions, time duration, and/or amount of data for collection, destination of collected data, intermediate processing, if any, and any other useful information (e.g., instance identifier, and the like). The controller 7802 may configure and operate the data collection system to perform the collection for the ODSV template and optionally return the system configuration to a previous configuration.

[0553] An example method of data collection for performing ODSV in an industrial environment includes automatically configuring local and remote data collection resources and collecting data from a number of sensors using the configured resources, where the number of sensors include a group of sensors that produce data that is required to perform the ODSV. In certain further embodiments, an example method further includes where the sensors are distributed throughout structural portions of an industrial machine in the industrial environment; where the sensors sense a range of system conditions including vibration, rotation, balance, and/or friction; where the automatically configuring is in response to a condition in the environment being detected outside of an acceptable range of condition values; where the condition is sensed by a sensor in a group of system sensors; where automatically configuring includes configuring a signal switching resource to concurrently connect a portion of the group of sensors to data collection resources; and/or where the signal switching resource is configured to maintain a connection between a reference sensor and the data collection resources throughout a period of collecting data from the sensors to perform ODSV.

[0554] An example method of data collection in an industrial environment includes configuring a data collection plan to collect data from a number of system sensors distributed throughout a machine in the industrial environment, the plan based on machine structural information and an indication of data needed to produce an ODSV of the machine; configuring data sensing, routing and collection resources in the environment based on the data collection plan; and collecting data based on the data collection plan. In certain further embodiments, an example method further includes: producing the ODSV; where the configuring data sensing, routing, and collection resources is in response to a condition in the environment being detected outside of an acceptable

range of condition values; where the condition is sensed by a sensor identified in the data collection plan; where configuring resources includes configuring a signal switching resource to concurrently connect the plurality of system sensors to data collection resources; and/or where the signal switching resource is configured to maintain a connection between a reference sensor and the data collection resources throughout a period of collecting data from the sensors to perform ODSV.

[0555] An example system for data collection in an industrial environment includes: a number of sensors disposed throughout the environment; multiplexer that connects signals from the plurality of sensors to data collection resources; and a processor for processing data collected from the number of sensors in response to the data collection template, where the processing results in an ODSV of a portion of a machine disposed in the environment. In certain further embodiments, an example system includes: where the ODSV collection template further identifies a condition in the environment on which performing data collection from the identified sensors is dependent; where the condition is sensed by a sensor identified in the ODSV data collection template; where the data collection template specified inputs of the multiplexer to concurrently connect to data collection resources; where the multiplexer is configured to maintain a connection between a reference sensor and the data collection resources throughout a period of collecting data from the sensors to perform ODSV; where the ODSV data collection template specifies data collection requirements for performing ODSV for looseness, soft joints, bending, and/or twisting of a portion of a machine in the industrial environment; and/or where the ODSV collection template specifies an order and timing of data collection from a plurality of identified sensors.

[0556] An example method of monitoring a mining conveyer for performing ODSV of the conveyer includes automatically configuring local and remote data collection resources and collecting data from a number of sensors disposed to sense the mining conveyor using the configured resources, wherein the plurality of sensors comprises a group of sensors that produce data that is required to perform the operational deflection shape visualization of a portion of the conveyor. An example method of monitoring a mining fan for performing ODSV of the fan includes automatically configuring local and remote data collection resources collecting data from a number of sensors disposed to sense the fan using the configured resources, and where the number of sensors include a group of sensors that produce data that is sufficient or required to perform ODSV of a portion of the fan.

[0557] In example embodiments, a system for data collection in an industrial environment may include a hierarchical multiplexer that facilitates successive multiplexing of input data channels according to a configurable hierarchy, such as a user configurable hierarchy. The system for data collection in an industrial environment may include the hierarchical multiplexer that facilitates successive multiplexing of a plurality of input data channels according to a configurable hierarchy. The hierarchy may be automatically configured by a controller based on an operational parameter in the industrial environment, such as a parameter of a machine in the industrial environment.

[0558] In example embodiments, a system for data collection in an industrial environment may include a plurality of sensors that may output data at different rates. The system may also include a multiplexer module that receives sensor outputs from a first portion of the plurality of sensors

with similar output rates into separate inputs of a first hierarchical multiplexer of the multiplexer module. The first hierarchical multiplexer of the multiplexer module may provide at least one multiplexed output of a portion of its inputs to a second hierarchical multiplexer that receives sensor outputs from a second portion of the plurality of sensors with similar output rates and that provides at least one multiplexed output of a portion of its inputs. In example embodiments, the output rates of the first set of sensors may be slower than the output rates of the second set of sensors. In example embodiments, data collection rate requirements of the first set of sensors may be lower than the data collection rate requirements of the second set of sensors. In example embodiments, the first hierarchical multiplexer output is a time-multiplexed combination of a portion of its inputs. In example embodiments, the second hierarchical multiplexer receives sensor signals with output rates that are similar to a rate of output of the first multiplexer, wherein the first multiplexer produces time-based multiplexing of the portion of its plurality of inputs.

[0559] In example embodiments, a system for data collection in an industrial environment may include a hierarchical multiplexer that is dynamically configured based on a data acquisition template. The hierarchical multiplexer may include a plurality of inputs and a plurality of outputs, wherein any input can be directed to any output in response to sensor output collection requirements of the template, and wherein a subset of the inputs can be multiplexed at a first switching rate and output to at least one of the plurality of outputs.

[0560] In example embodiments, a system for data collection in an industrial environment may include a plurality of sensors for sensing conditions of a machine in the environment, a hierarchical multiplexer, a plurality of analog-to-digital converters (ADCs), a processor, local storage, and an external interface. The system may use the processor to access a data acquisition template of parameters for data collection from a portion of the plurality of sensors, configure the hierarchical multiplexer, the ADCs and the local storage to facilitate data collection based on the defined parameters, and execute the data collection with the configured elements including storing a set of data collected from a portion of the plurality of sensors into the local storage. In example embodiments, the ADCs convert analog sensor data into a digital form that is compatible with the hierarchical multiplexer. In example embodiments, the processor monitors at least one signal generated by the sensors for a trigger condition and, upon detection of the trigger condition, responds by at least one of communicating an alert over the external interface and performing data acquisition according to a template that corresponds to the trigger condition.

[0561] In example embodiments, a system for data collection in an industrial environment may include a hierarchical multiplexer that may be configurable based on a data collection template of the environment. The multiplexer may support receiving a large number of data signals (e.g., from sensors in the environment) simultaneously. In example embodiments, all sensors for a portion of an industrial machine in the environment may be individually connected to inputs of a first stage of the multiplexer. The first stage of the multiplexer may provide a plurality of outputs that may feed into a second multiplexer stage. The second stage multiplexer may provide multiple outputs that feed into a third stage, and so on. Data collection templates for the environment may be configured for certain data collection sets, such as a set to determine temperature throughout a

machine or a set to determine vibration throughout a machine, and the like. Each template may identify a plurality of sensors in the environment from which data is to be collected, such as during a data collection event. When a template is presented to the hierarchical multiplexer, mapping of inputs to outputs for each multiplexing stage may be configured so that the required data is available at output(s) of a final multiplexing hierarchical stage for data collection. In an example, a data collection template to collect a set of data to determine temperature throughout a machine in the environment may identify many temperature sensors. The first stage multiplexer may respond to the template by selecting all of the available inputs that connect to temperature sensors. The data from these sensors may be multiplexed onto multiple inputs of a second stage sensor that may perform time-based multiplexing to produce a time-multiplexed output(s) of temperature data from a portion of the sensors. These outputs may be gathered by a data collector and de-multiplexed into individual sensor temperature readings.

[0562] In example embodiments, time-sensitive signals, such as triggers and the like, may connect to inputs that directly connect to a final multiplexer stage, thereby reducing any potential delay caused by routing through multiple multiplexing stages.

[0563] In example embodiments, a hierarchical multiplexer in a system for data collection in an industrial environment may comprise an array of relays, a programmable logic component, such as a CPLD, a field programmable gate array (FPGA), and the like.

[0564] In example embodiments, a system for data collection in an industrial environment that may include a hierarchical multiplexer for routing sensor outputs onto signal paths may be used with explosive systems in mining applications. Blast initiating and electronic blasting systems may be configured to provide computer assisted blasting systems. Ensuring that blasting occurs safely may involve effective sensing and analysis of a range of conditions. A system for data collection in an industrial environment may be deployed to sense and collect data associated with explosive systems, such as explosive systems used for mining. A data collection system can use a hierarchical multiplexer to capture data from explosive system installations automatically by aligning, for example, a deployment of the explosive system including its layout plans, integration, interconnectivity, cascading plan, and the like with the hierarchical multiplexer. An explosive system may be deployed with a form of hierarchy that starts with a primary initiator and follows detonation connections through successive layers of electronic blast control to sequenced detonation. Data collected from each of these layers of blast systems configuration may be associated with stages of a hierarchical multiplexer so that data collected from bulk explosive detonation can be captured in a hierarchy that corresponds to its blast control hierarchy.

[0565] In example embodiments, a system for data collection in an industrial environment that may include a hierarchical multiplexer for routing sensor outputs onto signal paths may be used with refinery blowers in oil and gas pipeline applications. Refinery blower applications include fired heater combustion air preheat systems and the like. Forced draft blowers may include a range of moving and moveable parts that may benefit from condition sensing and monitoring. Sensing may include detecting conditions of: couplings (e.g., temperature, rotational rate, and the like); motors (vibration, temperature, RPMs, torque, power usage, and the like); louver mechanics

(actuators, louvers, and the like); and plenums (flow rate, blockage, back pressure, and the like). A system for data collection in an industrial environment that uses a hierarchical multiplexer for routing signals from sensors and the like to data collectors may be configured to collect data from a refinery blower. In an example, a plurality of sensors may be deployed to sense air flow into, throughout, and out of a forced draft blower used in a refinery application, such as to preheat combustion air. Sensors may be grouped based on a frequency of a signal produced by sensors. Sensors that detect louver position and control may produce data at a lower rate than sensors that detect blower RPMs. Therefore, louver position and control sensor signals can be applied to a lower stage in a multiplexer hierarchy than the blower RPM sensors because data from louvers change less often than data from RPM sensors. A data collection system could switch among a plurality of louver sensors and still capture enough information to properly detect louver position. However, properly detecting blower RPM data may require greater bandwidth of connection between the blower RPM sensor and a data collector. A hierarchical multiplexer may enable capturing blower RPM data at a rate that is required for proper detection (perhaps by outputting the RPM sensor data for long durations of time), while switching among several louver sensor inputs and directing them onto (or through) an output that is different than the blower RPM output. Alternatively, the louver inputs may be time-multiplexed with the blower RPM data onto a single output that can be de-multiplexed by a data collector that is configured to determine when blower RPM data is being output and when louver position data is being output.

[0566] In example embodiments, a system for data collection in an industrial environment that may include a hierarchical multiplexer for routing sensor outputs onto signal paths may be used with pipeline-related compressors (e.g., reciprocating) in oil and gas pipeline applications. A typical use of a reciprocating compressor for pipeline application is production of compressed air for pipeline testing. A system for data collection in an industrial environment may apply a hierarchical multiplexer while collecting data from a pipeline testing-based reciprocating compressor. Data from sensors deployed along a portion of a pipeline being tested may be input to the lowest stage of the hierarchical multiplexer because these sensors may be periodically sampled prior to and during testing. However, the rate of sampling may be low relative to sensors that detect compressor operation, such as parts of the compressor that operate at higher frequencies, such as the reciprocating linkage, motor, and the like. The sensors that provide data at frequencies that enable reproduction of the detected motion may be input to higher stages in the hierarchical multiplexer. Time multiplexing among the pipeline sensors may provide for coverage of a large number of sensors while capturing events such as seal leakage and the like. However, time multiplexing among reciprocating linkage sensors may require output signal bandwidth that may exceed the bandwidth available for routing data from the multiplexer to a data collector. Therefore, in example embodiments, a plurality of pipeline sensors may be time-multiplexed onto a single multiplexer output and a compressor sensor detecting rapidly moving parts, such as the compressor motor, may be routed to separate outputs of the multiplexer.

[0567] Referring to Figure 39, a system for data collection in an industrial environment that uses a hierarchical multiplexer for routing sensor signals to data collectors is depicted. Outputs from a

plurality of sensors, such as sensors that monitor conditions that change with relatively low frequency (e.g., blower louver position sensors) may be input to a lowest hierarchical stage 8000 of a hierarchical multiplexer 8002 and routed to successively higher stages in the multiplexer, ultimately being output from the multiplexer, perhaps as a time-multiplexed signal comprising time-specific samples of each of the plurality of low frequency sensors. Outputs from a second plurality of sensors, such as sensors that monitor motor operation that may run at more than 1000 RPMs may be input to a higher hierarchical stage 8004 of the hierarchical multiplexer and routed to outputs that support the required bandwidth.

[0568] An example system for data collection in an industrial environment includes a controller for controlling data collection resources in the industrial environment and a hierarchical multiplexer that facilitates successive multiplexing of a number of input data channels according to a configurable hierarchy, wherein the hierarchy is automatically configured by the controller based on an operational parameter of a machine in the industrial environment. In certain further embodiments, an example system includes: where the operational parameter of the machine is identified in a data collection template; where the hierarchy is automatically configured in response to smart band data collection activation further including an analog-to-digital converter disposed between a source of the input data channels and the hierarchical multiplexer; and/or where the operational parameter of the machine comprises a trigger condition of at least one of the data channels. Another example system for data collection in an industrial environment includes a plurality of sensors and a multiplexer module that receives sensor outputs from a first portion of the sensors with similar output rates into separate inputs of a first hierarchical multiplexer that provides at least one multiplexed output of a portion of its inputs to a second hierarchical multiplexer, the second hierarchical multiplexer receiving sensor outputs from a second portion of the sensors and providing at least one multiplexed output of a portion of its inputs. In certain further embodiments, an example system includes: where the second portion of the sensors output data at rates that are higher than the output rates of the first portion of the sensors; where the first portion and the second portion of the sensors output data at different rates; where the first hierarchical multiplexer output is a time-multiplexed combination of a portion of its inputs; where the second multiplexer receives sensor signals with output rates that are similar to a rate of output of the first multiplexer; and/or where the first multiplexer produces time-based multiplexing of the portion of its inputs.

[0569] An example system for data collection in an industrial environment includes a number of sensors for sensing conditions of a machine in the environment a hierarchical multiplexer, a number of analog-to-digital converters, a controller, local storage, an external interface, where the system includes using the controller to access a data acquisition template that defines parameters for data collection from a portion of the sensors, to configure the hierarchical multiplexer, the ADCs, and the local storage to facilitate data collection based on the defined parameters, and to execute the data collection with the configured elements including storing a set of data collected from a portion of the sensors into the local storage. In certain further embodiments, an example system includes: where the ADCs convert analog sensor data into a digital form that is compatible

with the hierarchical multiplexer; where the processor monitors at least one signal generated by the sensors for a trigger condition and, upon detection of the trigger condition, responds by communicating an alert over the external interface and/or performing data acquisition according to a template that corresponds to the trigger condition; where the hierarchical multiplexer performs successive multiplexing of data received from the sensors according to a configurable hierarchy; where the hierarchy is automatically configured by the controller based on an operational parameter of a machine in the industrial environment; where the operational parameter of the machine is identified in a data collection template; where the hierarchy is automatically configured in response to smart band data collection activation; the system further including an ADC disposed between a source of the input data channels and the hierarchical multiplexer; where the operational parameter of the machine includes a trigger condition of at least one of the data channels; where the hierarchical multiplexer performs successive multiplexing of data received from the plurality of sensors according to a configurable hierarchy; and/or where the hierarchy is automatically configured by a controller based on a detected parameter of an industrial environment. Without limitation, n example system is configured for monitoring a mining explosive system, and includes a controller for controlling data collection resources associated with the explosive system, and a hierarchical multiplexer that facilitates successive multiplexing of a number of input data channels according to a configurable hierarchy, where the hierarchy is automatically configured by the controller based on a configuration of the explosive system. Without limitation, an example system is configured for monitoring a refinery blower in an oil and gas pipeline applications, and includes a controller for controlling data collection resources associated with the refinery blower, and a hierarchical multiplexer that facilitates successive multiplexing of a number of input data channels according to a configurable hierarchy, where the hierarchy is automatically configured by the controller based on a configuration of the refinery blower. Without limitation, an example system is configured for monitoring a reciprocating compressor in an oil and gas pipeline applications comprising, and includes controller for controlling data collection resources associated with the reciprocating compressor, and a hierarchical multiplexer that facilitates successive multiplexing of a number of input data channels according to a configurable hierarchy, where the hierarchy is automatically configured by the controller based on a configuration of the reciprocating compressor.

[0570] In example embodiments, a system for data collection in an industrial environment may include an ultrasonic sensor disposed to capture ultrasonic conditions of an element of in the environment. The system may be configured to collect data representing the captured ultrasonic condition in a computer memory, on which a processor may execute an ultrasonic analysis algorithm. In example embodiments, the sensed element may be one of a moving element, a rotating element, a structural element, and the like. In example embodiments, the data may be streamed to the computer memory. In example embodiments, the data may be continuously streamed. In example embodiments, the data may be streamed for a duration of time, such as an ultrasonic condition sampling duration. In example embodiments, the system may also include a data routing infrastructure that facilitates routing the streaming data from the ultrasonic sensor to

a plurality of destinations including local and remote destinations. The routing infrastructure may include a hierarchical multiplexer that is adapted to route the streaming data and data from at least one other sensor to a destination.

[0571] In example embodiments, ultrasonic monitoring in an industrial environment may be performed by a system for data collection as described herein on rotating elements (e.g., motor shafts and the like), bearings, fittings, couplings, housings, load bearing elements, and the like. The ultrasonic data may be used for pattern recognition, state determination, time-series analysis, and the like, any of which may be performed by computing resources of the industrial environment, which may include local computing resources (e.g., resources located within the environment and/or within a machine in the environment, and the like) and remote computing resources (e.g., cloud-based computing resources, and the like).

[0572] In example embodiments, ultrasonic monitoring in an industrial environment by a system for data collection may be activated in response to a trigger (e.g., a signal from a motor indicating the motor is operational, and the like), a measure of time (e.g., an amount of time since the most recent monitoring activity, a time of day, a time relative to a trigger, an amount of time until a future event, such as machine shutdown, and the like), an external event (e.g., lightning strike, and the like). The ultrasonic monitoring may be activated in response to implementation of a smart band data collection activity. The ultrasonic monitoring may be activated in response to a data collection template being applied in the industrial environment. The data collection template may be configured based on analysis of prior vibration-caused failures that may be applicable to the monitored element, machine, environment, and the like. Because continuous monitoring of ultrasonic data may require dedicating data routing resources in the industrial environment for extended periods of time, a data collection template for continuous ultrasonic monitoring may be configured with data routing and resource utilization setup information that a controller of a data collection system may use to setup the resources to accommodate continuous ultrasonic monitoring. In an example, a data multiplexer may be configured to dedicate a portion of its outputs to the ultrasonic data for a duration of time specified in the template.

[0573] In example embodiments, a system for data collection in an industrial environment may perform continuous ultrasonic monitoring. The system may also include processing of the ultrasonic data by a local processor located proximal to the vibration monitoring sensor or device(s). Depending on the computing capabilities of the local processor, functions such as peak detection may be performed. A programmable logic component may provide sufficient computing capabilities to perform peak detection. Processing of the ultrasonic data (local or remote) may provide feedback to a controller associated with the element(s) being monitored. The feedback may be used in a control loop to potentially adjust an operating condition, such as rotational speed, and the like, in an attempt to reduce or at least contain potential negative impact suggested by the ultrasonic data analysis.

[0574] In example embodiments, a system for data collection in an industrial environment may perform ultrasonic monitoring, and in particular, continuous ultrasonic monitoring. The ultrasonic monitoring data may be combined with multi-dimensional models of an element or machine being

monitored to produce a visualization of the ultrasonic data. In example embodiments, an image, set of images, video, and the like may be produced that correlates in time with the sensed ultrasonic data. In example embodiments, image recognition and/or analysis may be applied to ultrasonic visualizations to further facilitate determining the severity of a condition detected by the ultrasonic monitoring. The image analysis algorithms may be trained to detect normal and out of bounds conditions. Data from load sensors may be combined with ultrasonic data to facilitate testing materials and systems.

[0575] In example embodiments, a system for data collection in an industrial environment may perform ultrasonic monitoring of a pipeline in an oil and gas pipeline application. Flows of petroleum through pipelines can create vibration and other mechanical effects that may contribute to structural changes in a liner of the pipeline, support members, flow boosters, regulators, diverters, and the like. Performing continuous ultrasonic monitoring of key elements in a pipeline may facilitate detecting early changes in material, such as joint fracturing, and the like, that may lead to failure. A system for data collection in an industrial environment may be configured with ultrasonic sensing devices that may be connected through signal data routing resources, such as crosspoint switches, multiplexers, and the like, to data collection and analysis nodes at which the collected ultrasonic data can be collected and analyzed. In example embodiments, a data collection system may include a controller that may reference a data collection plan or template that includes information to facilitate configuring the data sampling, routing, and collection resources of the system to accommodate collecting ultrasonic sample data from a plurality of elements along the pipeline. The template may indicate a sequence for collecting ultrasonic data from a plurality of ultrasonic sensors and the controller may configure a multiplexer to route ultrasonic sensor data from a specified ultrasonic sensor to a destination, such as a data storage controller, analysis processor and the like, for a duration specified in the template. The controller may detect a sequence of collection in the template, or a sequence of templates to access, and respond to each template in the detected sequence, adjusting the multiplexer and the like to route the sensor data specified in each template to a collector.

[0576] In example embodiments, a system for data collection in an industrial environment may perform ultrasonic monitoring of compressors in a power generation application. Compressors include several critical rotating elements (e.g., shaft, motor, and the like), rotational support elements (e.g., bearings, couplings, and the like), and the like. A system for data collection configured to facilitate sensing, routing, collection and analysis of ultrasonic data in a power generation application may receive ultrasonic sensor data from a plurality of ultrasonic sensors. Based on a configuration setup template, such as a template for collecting continuous ultrasonic data from one or more ultrasonic sensor devices, a controller may configure resources of the data collection system to facilitate delivery of the ultrasonic data over one or more signal data lines from the sensor(s) at least to data collectors that may be locally or remotely accessible. In example embodiments, a template may indicate that ultrasonic data for a main shaft should be retrieved continuously for one minute, and then ultrasonic data for a secondary shaft should be retrieved for another minute, followed by ultrasonic data for a housing of the compressor. The controller may

configure a multiplexer that receives the ultrasonic data for each of these sensors to route the data from each sensor in order by configuring a control set that initially directs the inputs from the main shaft ultrasonic sensors through the multiplexer until the time or other measure of data being forwarded is reached. The controller could switch the multiplexer to route the additional ultrasonic data as required to satisfy the second template requirements. The controller may continue adjusting the data collection system resources along the way until all of the ultrasonic monitoring data collection templates are satisfied.

[0577] In example embodiments, a system for data collection in an industrial environment may perform ultrasonic monitoring of wind turbine gearboxes in a wind energy generation application. Gearboxes in wind turbines may experience a high degree of resistance in operation, due in part to the changing nature of wind, which may cause moving parts, such as the gear planes, hydraulic fluid pumps, regulators, and the like, to prematurely fail. A system for data collection in an industrial environment may be configured with ultrasonic sensors that capture information that may lead to early detection of potential failure modes of these high-strain elements. To ensure that ultrasonic data may be effectively acquired from several different ultrasonic sensors with sufficient coverage to facilitate producing an actionable ultrasonic imaging assessment, the system may be configured specifically to deliver sufficient data at a relatively high rate from one or more of the sensors. Routing channel(s) may be dedicated to transferring ultrasonic sensing data for a duration of time that may be specified in an ultrasonic data collection plan or template. To accomplish this, a controller, such as a programmable logic component, may configure a portion of a crosspoint switch and data collectors to deliver ultrasonic data from a first set of ultrasonic sensors (e.g., those that sense hydraulic fluid flow control elements) to a plurality of data collectors. Another portion of the crosspoint switch may be configured to route additional sensor data that may be useful for evaluating the ultrasonic data (e.g., motor on/off state, thermal condition of sensed parts, and the like) on other data channels to data collectors where the data can be combined and analyzed. The controller may reconfigure the data routing resources to enable collecting ultrasonic data from other elements based on a corresponding data collection template.

[0578] Referring to Figure 40, a system for data collection in an industrial environment may include one or more ultrasonic sensors 8050 that may connect to a data collection and routing system 8052 that may be configured by a controller 8054 based on an ultrasonic sensor-specific data collection template 8056 that may be provided to the controller 8054 by an ultrasonic data analysis facility 8058. The controller 8054 may configure resources of the data collection system 8052 and monitor the data collection for a duration of time based on the requirements for data collection in the template 8056.

[0579] In example embodiments, information about the health or other status or state information of or regarding a component or piece of industrial equipment may be obtained by monitoring the condition of various components throughout a process. Monitoring may include monitoring the amplitude of a sensor signal measuring attributes such as temperature, humidity, acceleration, displacement, and the like. An embodiment of a data monitoring device 8100 is shown in Figure 41 and may include a plurality of sensors 8106 communicatively coupled to a controller 8102. The

controller 8102 may include a data acquisition circuit 8104, a data analysis circuit 8108, a MUX control circuit 8114, and a response circuit 8110. The data acquisition circuit 8104 may include a MUX 8112 where the inputs correspond to a subset of the detection values. The MUX control circuit 8114 may be structured to provide adaptive scheduling of the logical control of the MUX and the correspondence of MUX input and detected values based on a subset of the plurality of detection values and/or a command from the response circuit 8110 and/or the output of the data analysis circuit 8104. The data analysis circuit 8108 may comprise one or more of a peak detection circuit, a phase differential circuit, a PLL circuit, a bandpass filter circuit, a frequency transformation circuit, a frequency analysis circuit, a torsional analysis circuit, a bearing analysis circuit, an overload detection circuit, a sensor fault detection circuit, a vibrational resonance circuit for the identification of unfavorable interaction among machines or components, a distortion identification circuit for the identification of unfavorable distortions such as deflections shapes upon operation, overloading of weight, excessive forces, stress and strain-based effects, and the like. The data analysis circuit 8108 may output a component health status as a result of the analysis.

[0580] The data analysis circuit 8108 may determine a state, condition, or status of a component, part, sub-system, or the like of a machine, device, system or item of equipment (collectively referred to herein as a component health status) based on a maximum value of a MUX output for a given input or a rate of change of the value of a MUX output for a given input. The data analysis circuit 8108 may determine a component health status based on a time integration of the value of a MUX for a given input. The data analysis circuit 8108 may determine a component health status based on phase differential of MUX output relative to an on-board time or another sensor. The data analysis circuit 8108 may determine a component health status based on a relationship of value, phase, phase differential, and rate of change for MUX outputs corresponding to one or more input detection values. The data analysis circuit 8108 may determine a component health status based on process stage or component specification or component anticipated state.

[0581] The multiplexer control circuit 8114 may adapt the scheduling of the logical control of the multiplexer based on a component health status, an anticipated component health status, the type of component, the type of equipment being measured, an anticipated state of the equipment, a process stage (different parameters/sensor values) may be important at different stages in a process. The multiplexer control circuit 8114 may adapt the scheduling of the logical control of the multiplexer based on a sequence selected by a user or a remote monitoring application, or on the basis of a user request for a specific value. The multiplexer control circuit 8114 may adapt the scheduling of the logical control of the multiplexer based on the basis of a storage profile or plan (such as based on type and availability of storage elements and parameters as described elsewhere in this disclosure and in the documents incorporated herein by reference), network conditions or availability (also as described elsewhere in this disclosure and in the documents incorporated herein by reference), or value or cost of component or equipment.

[0582] The plurality of sensors 8106 may be wired to ports on the data acquisition circuit 8104. The plurality of sensors 8106 may be wirelessly connected to the data acquisition circuit 8104. The data acquisition circuit 8104 may be able to access detection values corresponding to the output of

at least one of the plurality of sensors 8106 where the sensors 8106 may be capturing data on different operational aspects of a piece of equipment or an operating component.

[0583] The selection of the plurality of sensors 8106 for a data monitoring device 8100 designed for a specific component or piece of equipment may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, resolution desired at various positions in a process or plant, reliability of the sensors, and the like. The impact of a failure, time response of a failure (e.g., warning time and/or off-nominal modes occurring before failure), likelihood of failure, and/or sensitivity required, and/or difficulty to detect failure conditions may drive the extent to which a component or piece of equipment is monitored with more sensors, and/or higher capability sensors being dedicated to systems where unexpected or undetected failure would be costly or have severe consequences.

[0584] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating, and the like, sensors 8106 may comprise one or more of, without limitation, a vibration sensor, a thermometer, a hygrometer, a voltage sensor and/or a current sensor (for the component and/or other sensors measuring the component), an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition, and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, a thermal imager, an acoustic wave sensor, a displacement sensor, a turbidity meter, a viscosity meter, an axial load sensor, a radial load sensor, a tri-axial sensor, an accelerometer, a speedometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an optical (laser) particle counter, an ultrasonic sensor, an acoustical sensor, a heat flux sensor, a galvanic sensor, a magnetometer, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference.

[0585] The sensors 8106 may provide a stream of data over time that has a phase component, such as relating to acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 8106 may provide a stream of data that is not conventionally phase-based, such as temperature, humidity, load, and the like. The sensors 8106 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0586] The sensors 8106 may monitor components such as bearings, sets of bearings, motors, driveshafts, pistons, pumps, conveyors, vibrating conveyors, compressors, drills, and the like in vehicles, oil and gas equipment in the field, in assembly line components, and the like.

[0587] In example embodiments, as illustrated in Figure 41, the sensors 8106 may be part of the data monitoring device 8100, referred to herein in some cases as a data collector, which in some cases may comprise a mobile or portable data collector. In example embodiments, as illustrated in Figures 42 and 43, one or more external sensors 8126, which are not explicitly part of a monitoring device 8120 but rather are new, previously attached to or integrated into the equipment or

component, may be opportunistically connected to, or accessed by the monitoring device 8120. The monitoring device 8120 may include a controller 8122. The controller 8122 may include a data acquisition circuit 8104, a data analysis circuit 8108, a MUX control circuit 8114, and a response circuit 8110. The data acquisition circuit 8104 may comprise a MUX 8112 where the inputs correspond to a subset of the detection values. The MUX control circuit 8114 may be structured to provide the logical control of the MUX and the correspondence of MUX input and detected values based on a subset of the plurality of detection values and/or a command from the response circuit 8110 and/or the output of the data analysis circuit 8108. The data analysis circuit 8108 may comprise one or more of a peak detection circuit, a phase differential circuit, a PLL circuit, a bandpass filter circuit, a frequency transformation circuit, a frequency analysis circuit, a torsional analysis circuit, a bearing analysis circuit, an overload detection circuit, vibrational resonance circuit for the identification of unfavorable interaction among machines or components, a distortion identification circuit for the identification of unfavorable distortions such as deflections shapes upon operation ,stress and strain-based effects, and the like.

[0588] The one or more external sensors 8126 may be directly connected to the one or more input ports 8128 on the data acquisition circuit 8104 of the controller 8122 or may be accessed by the data acquisition circuit 8104 wirelessly, such as by a reader, interrogator, or other wireless connection, such as over a short-distance wireless protocol. In example embodiments, as shown in Figure 43, a data acquisition circuit 8104 may further comprise a wireless communication circuit 8130. The data acquisition circuit 8104 may use the wireless communication circuit 8130 to access detection values corresponding to the one or more external sensors 8126 wirelessly or via a separate source or some combination of these methods.

[0589] In example embodiments, as illustrated in Figure 44, the controller 8134 may further comprise a data storage circuit 8136. The data storage circuit 8136 may be structured to store one or more of sensor specifications, component specifications, anticipated state information, detected values, multiplexer output, component models, and the like. The data storage circuit 8136 may provide specifications and anticipated state information to the data analysis circuit 8108.

[0590] In example embodiments, the response circuit 8110 may initiate a variety of actions based on the sensor status provided by the data analysis circuit 8108. The response circuit 8110 may adjust a sensor scaling value (e.g., from 100 mV/gram to 10 mV/gram). The response circuit 8110 may select an alternate sensor from a plurality available. The response circuit 8110 may acquire data from a plurality of sensors of different ranges. The response circuit 8110 may recommend an alternate sensor. The response circuit 8110 may issue an alarm or an alert.

[0591] In example embodiments, the response circuit 8110 may cause the data acquisition circuit 8104 to enable or disable the processing of detection values corresponding to certain sensors based on the component status. This may include switching to sensors having different response rates, sensitivity, ranges, and the like; accessing new sensors or types of sensors, accessing data from multiple sensors, and the like. Switching may be undertaken based on a model, a set of rules, or the like. In example embodiments, switching may be under control of a machine learning system, such that switching is controlled based on one or more metrics of success, combined with input

data, over a set of trials, which may occur under supervision of a human supervisor or under control of an automated system. Switching may involve switching from one input port to another (such as to switch from one sensor to another). Switching may involve altering the multiplexing of data, such as combining different streams under different circumstances. Switching may involve activating a system to obtain additional data, such as moving a mobile system (such as a robotic or drone system), to a location where different or additional data is available, such as positioning an image sensor for a different view or positioning a sonar sensor for a different direction of collection, or to a location where different sensors can be accessed, such as moving a collector to connect up to a sensor at a location in an environment by a wired or wireless connection. This switching may be implemented by directing changes to the multiplexer (MUX) control circuit 8114.

[0592] In example embodiments, the response circuit 8110 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 8110 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0593] In example embodiments, the response circuit 8110 may recommend maintenance at an upcoming process stop or initiate a maintenance call where the maintenance may include the replacement of the sensor with the same or an alternate type of sensor having a different response rate, sensitivity, range, and the like. In example embodiments, the response circuit 8110 may implement or recommend process changes – for example to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but is still at least partially operational, to change the operating speed of a component (such as to put it in a lower-demand mode), to initiate amelioration of an issue (such as to signal for additional lubrication of a roller bearing set, or to signal for an alignment process for a system that is out of balance), and the like.

[0594] In example embodiments, the data analysis circuit 8108 and/or the response circuit 8110 may periodically store certain detection values and/or the output of the multiplexers and/or the data corresponding to the logic control of the MUX in the data storage circuit 8136 to enable the tracking of component performance over time. In example embodiments, based on sensor status, as described elsewhere herein, recently measured sensor data and related operating conditions such as RPMs, component loads, temperatures, pressures, vibrations, or other sensor data of the types described throughout this disclosure in the data storage circuit 8136 enable the backing out of overloaded/failed sensor data. The signal evaluation circuit 8108 may store data at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0595] In example embodiments, as shown in Figures 45, 46, 47, and 48, a data monitoring system 8138 may include at least one data monitoring device 8140. The at least one data monitoring device 8140 may include sensors 8106 and a controller 8142 comprising a data acquisition circuit 8104,

a data analysis circuit 8108, a data storage circuit 8136, and a communication circuit 8146 to allow data and analysis to be transmitted to a monitoring application 8150 on a remote server 8148. The signal evaluation circuit 8108 may include at least an overload detection circuit (e.g., reference Figures 91 and 92) and/or a sensor fault detection circuit (e.g., reference Figures 91 and 92). The signal evaluation circuit 8108 may periodically share data with the communication circuit 8146 for transmittal to the remote server 8148 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 8150. Based on the sensor status, the signal evaluation circuit 8108 and/or response circuit 8110 may share data with the communication circuit 8146 for transmittal to the remote server 8148 based on the fit of data relative to one or more criteria. Data may include recent sensor data and additional data such as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The signal evaluation circuit 8108 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server.

[0596] In example embodiments, as shown in Figure 45, the communication circuit 8146 may communicate data directly to a remote server 8148. In example embodiments, as shown in Figure 46, the communication circuit 8146 may communicate data to an intermediate computer 8152 which may include a processor 8154 running an operating system 8156 and a data storage circuit 8158.

[0597] In embodiments as illustrated in Figures 47 and 48, a data collection system 8160 may have a plurality of monitoring devices 8144 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment, (both the same and different types of equipment) in the same facility, as well as collecting data from monitoring devices in multiple facilities. A monitoring application 8150 on a remote server 8148 may receive and store one or more of detection values, timing signals, and data coming from a plurality of the various monitoring devices 8144.

[0598] In example embodiments, as shown in Figure 47, the communication circuit 8146 may communicate data directly to a remote server 8148. In example embodiments, as shown in Figure 48, the communication circuit 8146 may communicate data to an intermediate computer 8152 which may include a processor 8154 running an operating system 8156 and a data storage circuit 8158. There may be an individual intermediate computer 8152 associated with each monitoring device 8140 or an individual intermediate computer 8152 may be associated with a plurality of monitoring devices 8144 where the intermediate computer 8152 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 8148. Communication to the remote server 8148 may be streaming, batch (e.g., when a connection is available), or opportunistic.

[0599] The monitoring application 8150 may select subsets of the detection values to be jointly analyzed. Subsets for analysis may be selected based on a single type of sensor, component, or a single type of equipment in which a component is operating. Subsets for analysis may be selected or grouped based on common operating conditions such as size of load, operational condition (e.g., intermittent or continuous), operating speed or tachometer output, common ambient environmental

conditions such as humidity, temperature, air or fluid particulate, and the like. Subsets for analysis may be selected based on the effects of other nearby equipment such as nearby machines rotating at similar frequencies, nearby equipment producing electromagnetic fields, nearby equipment producing heat, nearby equipment inducing movement or vibration, nearby equipment emitting vapors, chemicals or particulates, or other potentially interfering or intervening effects.

[0600] In example embodiments, the monitoring application 8150 may analyze the selected subset. In an example, data from a single sensor may be analyzed over different time periods such as one operating cycle, several operating cycles, a month, a year, the life of the component, or the like. Data from multiple sensors of a common type measuring a common component type may also be analyzed over different time periods. Trends in the data such as changing rates of change associated with start-up or different points in the process may be identified. Correlation of trends and values for different sensors may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected sensor performance. This information may be transmitted back to the monitoring device to update sensor models, sensor selection, sensor range, sensor scaling, sensor sampling frequency, types of data collected, and the like, and be analyzed locally or to influence the design of future monitoring devices.

[0601] In example embodiments, the monitoring application 8150 may have access to equipment specifications, equipment geometry, component specifications, component materials, anticipated state information for a plurality of sensors, operational history, historical detection values, sensor life models, and the like for use analyzing the selected subset using rule-based or model-based analysis. The monitoring application 8150 may provide recommendations regarding sensor selection, additional data to collect, data to store with sensor data, and the like. The monitoring application 8150 may provide recommendations regarding scheduling repairs and/or maintenance. The monitoring application 8150 may provide recommendations regarding replacing a sensor. The replacement sensor may match the sensor being replaced or the replacement sensor may have a different range, sensitivity, sampling frequency, and the like.

[0602] In example embodiments, the monitoring application 8150 may include a remote learning circuit structured to analyze sensor status data (e.g., sensor overload or sensor failure) together with data from other sensors, failure data on components being monitored, equipment being monitored, output being produced, and the like. The remote learning system may identify correlations between sensor overload and data from other sensors.

[0603] An example monitoring system for data collection in an industrial environment includes a data acquisition circuit that interprets a number of detection values, each of the detection values corresponding to input received from at least one of a number of input sensors, a MUX having inputs corresponding to a subset of the detection values, a MUX control circuit that interprets a subset of the number of detection values and provides the logical control of the MUX and the correspondence of MUX input and detected values as a result, where the logic control of the MUX includes adaptive scheduling of the select lines, a data analysis circuit that receives an output from the MUX and data corresponding to the logic control of the MUX resulting in a component health status, an analysis response circuit that performs an operation in response to the component health

status, where the number of sensors includes at least two sensors such as a temperature sensor, a load sensor, a vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor, and/or a tachometer. In certain further embodiments, an example system includes: where at least one of the number of detection values may correspond to a fusion of two or more input sensors representing a virtual sensor; where the system further includes a data storage circuit that stores at least one of component specifications and anticipated component state information and buffers a subset of the number of detection values for a predetermined length of time; where the system further includes a data storage circuit that stores at least one of a component specification and anticipated component state information and buffers the output of the MUX and data corresponding to the logic control of the MUX for a predetermined length of time; where the data analysis circuit includes a peak detection circuit, a phase detection circuit, a bandpass filter circuit, a frequency transformation circuit, a frequency analysis circuit, a PLL circuit, a torsional analysis circuit, and/or a bearing analysis circuit; where operation further includes storing additional data in the data storage circuit; where the operation includes at least one of enabling or disabling one or more portions of the MUX circuit; and/or where the operation includes causing the MUX control circuit to alter the logical control of the MUX and the correspondence of MUX input and detected values. In certain example embodiments, the system includes at least two multiplexers; control of the correspondence of the multiplexer input and the detected values further includes controlling the connection of the output of a first multiplexer to an input of a second multiplexer; control of the correspondence of the multiplexer input and the detected values further comprises powering down at least a portion of one of the at least two multiplexers; and/or control of the correspondence of MUX input and detected values includes adaptive scheduling of the select lines. In certain example embodiments, a data response circuit analyzes the stream of data from one or both MUXes, and recommends an action in response to the analysis.

[0604] An example testing system includes the testing system in communication with a number of analog and digital input sensors, a monitoring device including a data acquisition circuit that interprets a number of detection values, each of the number of detection values corresponding to at least one of the input sensors, a MUX having inputs corresponding to a subset of the detection values, a MUX control circuit that interprets a subset of the number of detection values and provides the logical control of the MUX and control of the correspondence of MUX input and detected values as a result, where the logic control of the MUX includes adaptive scheduling of the select lines, and a user interface enabled to accept scheduling input for select lines and display output of MUX and select line data.

[0605] In example embodiments, information about the health or other status or state information of or regarding a component or piece of industrial equipment may be obtained by looking at both the amplitude and phase or timing of data signals relative to related data signals, timers, reference signals or data measurements. An embodiment of a data monitoring device 8500 is shown in Figure 49 and may include a plurality of sensors 8506 communicatively coupled to a controller 8502. The controller 8502 may include a data acquisition circuit 8504, a signal evaluation circuit 8508 and a

response circuit 8510. The plurality of sensors 8506 may be wired to ports on the data acquisition circuit 8504 or wirelessly in communication with the data acquisition circuit 8504. The plurality of sensors 8506 may be wirelessly connected to the data acquisition circuit 8504. The data acquisition circuit 8504 may be able to access detection values corresponding to the output of at least one of the plurality of sensors 8506 where the sensors 8506 may be capturing data on different operational aspects of a piece of equipment or an operating component.

[0606] The selection of the plurality of sensors 8506 for a data monitoring device 8500 designed for a specific component or piece of equipment may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, reliability of the sensors, and the like. The impact of failure may drive the extent to which a component or piece of equipment is monitored with more sensors and/or higher capability sensors being dedicated to systems where unexpected or undetected failure would be costly or have severe consequences.

[0607] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating and the like, sensors 8506 may comprise one or more of, without limitation, a vibration sensor, a thermometer, a hygrometer, a voltage sensor, a current sensor, an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, an acoustic wave sensor, a displacement sensor, a turbidity meter, a viscosity meter, a load sensor, a tri-axial sensor, an accelerometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an acoustical sensor, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference.

[0608] The sensors 8506 may provide a stream of data over time that has a phase component, such as relating to acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 8506 may provide a stream of data that is not conventionally phase-based, such as temperature, humidity, load, and the like. The sensors 8506 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0609] In example embodiments, as illustrated in Figure 49, the sensors 8506 may be part of the data monitoring device 8500, referred to herein in some cases as a data collector, which in some cases may comprise a mobile or portable data collector. In example embodiments, as illustrated in Figures 50 and 51, sensors 8518, either new or previously attached to or integrated into the equipment or component, may be opportunistically connected to or accessed by a monitoring device 8512. The sensors 8518 may be directly connected to input ports 8520 on the data acquisition circuit 8516 of a controller 8514 or may be accessed by the data acquisition circuit 8516 wirelessly, such as by a reader, interrogator, or other wireless connection, such as over a short-distance wireless protocol. In example embodiments, a data acquisition circuit 8516 may access detection values corresponding to the sensors 8518 wirelessly or via a separate source or

some combination of these methods. In example embodiments, the data acquisition circuit 8504 may include a wireless communications circuit 8522 able to wirelessly receive data opportunistically from sensors 8518 in the vicinity and route the data to the input ports 8520 on the data acquisition circuit 8516.

[0610] In example embodiments, as illustrated in Figures 52 and 53, the signal evaluation circuit 8508 may then process the detection values to obtain information about the component or piece of equipment being monitored. Information extracted by the signal evaluation circuit 8508 may comprise rotational speed, vibrational data including amplitudes, frequencies, phase, and/or acoustical data, and/or non-phase sensor data such as temperature, humidity, image data, and the like.

[0611] The signal evaluation circuit 8508 may include one or more components such as a phase detection circuit 8528 to determine a phase difference between two time-based signals, a phase lock loop circuit 8530 to adjust the relative phase of a signal such that it is aligned with a second signal, timer or reference signal, and/or a band pass filter circuit 8532 which may be used to separate out signals occurring at different frequencies. An example band pass filter circuit 8532 includes any filtering operations understood in the art, including at least a low-pass filter, a high-pass filter, and/or a band pass filter – for example to exclude or reduce frequencies that are not of interest for a particular determination, and/or to enhance the signal for frequencies of interest. Additionally, or alternatively, a band pass filter circuit 8532 includes one or more notch filters or other filtering mechanism to narrow ranges of frequencies (e.g., frequencies from a known source of noise). This may be used to filter out dominant frequency signals such as the overall rotation, and may help enable the evaluation of low amplitude signals at frequencies associated with torsion, bearing failure and the like.

[0612] In example embodiments, understanding the relative differences may be enabled by a phase detection circuit 8528 to determine a phase difference between two signals. It may be of value to understand a relative phase offset, if any, between signals such as when a periodic vibration occurs relative to a relative rotation of a piece of equipment. In example embodiments, there may be value in understanding where in a cycle shaft vibrations occur relative to a motor control input to better balance the control of the motor. This may be particularly true for systems and components that are operating at relatively slow RPMs. Understanding of the phase difference between two signals or between those signals and a timer may enable establishing a relationship between a signal value and where it occurs in a process or rotation. Understanding relative phase differences may help in evaluating the relationship between different components of a system such as in the creation of a vibrational model for an Operational Deflection Shape (ODS).

[0613] The signal evaluation circuit 8544 may perform frequency analysis using techniques such as a digital Fast Fourier transform (FFT), Laplace transform, Z-transform, wavelet transform, other frequency domain transform, or other digital or analog signal analysis techniques, including, without limitation, complex analysis, including complex phase evolution analysis. An overall rotational speed or tachometer may be derived from data from sensors such as rotational velocity meters, accelerometers, displacement meters and the like. Additional frequencies of interest may

also be identified. These may include frequencies near the overall rotational speed as well as frequencies higher than that of the rotational speed. These may include frequencies that are nonsynchronous with an overall rotational speed. Signals observed at frequencies that are multiples of the rotational speed may be due to bearing induced vibrations or other behaviors or situations involving bearings. In some instances, these frequencies may be in the range of one times the rotational speed, two times the rotational speed, three times the rotational speed, and the like, up to 3.15 to 15 times the rotational speed, or higher. In example embodiments, the signal evaluation circuit 8544 may select RC components for a band pass filter circuit 8532 based on overall rotational speed to create a band pass filter circuit 8532 to remove signals at expected frequencies such as the overall rotational speed, to facilitate identification of small amplitude signals at other frequencies. In example embodiments, variable components may be selected, such that adjustments may be made in keeping with changes in the rotational speed, so that the band pass filter may be a variable band pass filter. This may occur under control of automatically self-adjusting circuit elements, or under control of a processor, including automated control based on a model of the circuit behavior, where a rotational speed indicator or other data is provided as a basis for control.

[0614] In example embodiments, rather than performing frequency analysis, the signal evaluation circuit 8544 may utilize the time-based detection values to perform transitory signal analysis. These may include identifying abrupt changes in signal amplitude including changes where the change in amplitude exceeds a predetermined value or exists for a certain duration. In example embodiments, the time-based sensor data may be aligned with a timer or reference signal allowing the time-based sensor data to be aligned with, for example, a time or location in a cycle. Additional processing to look at frequency changes over time may include the use of Short-Time Fourier Transforms (STFT) or a wavelet transform.

[0615] In example embodiments, frequency-based techniques and time-based techniques may be combined, such as using time-based techniques to determine discrete time periods during which given operational modes or states are occurring and using frequency-based techniques to determine behavior within one or more of the discrete time periods.

[0616] In example embodiments, the signal evaluation circuit may utilize demodulation techniques for signals obtained from equipment running at slow speeds such as paper and pulp machines, mining equipment, and the like. A signal evaluation circuit employing a demodulation technique may comprise a band-pass filter circuit, a rectifier circuit, and/or a low pass circuit prior to transforming the data to the frequency domain.

[0617] The response circuit 8510 8710 may further comprise evaluating the results of the signal evaluation circuit 8508 8544 and, based on certain criteria, initiating an action. Criteria may include a predetermined maximum or minimum value for a detection value from a specific sensor, a value of a sensor's corresponding detection value over time, a change in value, a rate of change in value, and/or an accumulated value (e.g., a time spent above/below a threshold value, a weighted time spent above/below one or more threshold values, and/or an area of the detected value above/below one or more threshold values). The criteria may include a sensor's detection values at certain frequencies or phases where the frequencies or phases may be based on the equipment geometry,

equipment control schemes, system input, historical data, current operating conditions, and/or an anticipated response. The criteria may comprise combinations of data from different sensors such as relative values, relative changes in value, relative rates of change in value, relative values over time, and the like. The relative criteria may change with other data or information such as process stage, type of product being processed, type of equipment, ambient temperature and humidity, external vibrations from other equipment, and the like. The relative criteria may include level of synchronicity with an overall rotational speed, such as to differentiate between vibration induced by bearings and vibrations resulting from the equipment design. In example embodiments, the criteria may be reflected in one or more calculated statistics or metrics (including ones generated by further calculations on multiple criteria or statistics), which in turn may be used for processing (such as on board a data collector or by an external system), such as to be provided as an input to one or more of the machine learning capabilities described in this disclosure, to a control system (which may be an on-board data collector or remote, such as to control selection of data inputs, multiplexing of sensor data, storage, or the like), or as a data element that is an input to another system, such as a data stream or data package that may be available to a data marketplace, a SCADA system, a remote control system, a maintenance system, an analytic system, or other system.

[0618] In an illustrative and non-limiting example, an alert may be issued if the vibrational amplitude and/or frequency exceeds a predetermined maximum value, if there is a change or rate of change that exceeds a predetermined acceptable range, and/or if an accumulated value based on vibrational amplitude and/or frequency exceeds a threshold. Certain embodiments are described herein as detected values exceeding thresholds or predetermined values, but detected values may also fall below thresholds or predetermined values -- for example where an amount of change in the detected value is expected to occur, but detected values indicate that the change may not have occurred. For example, and without limitation, vibrational data may indicate system agitation levels, properly operating equipment, or the like, and vibrational data below amplitude and/or frequency thresholds may be an indication of a process that is not operating according to expectations. Except where the context clearly indicates otherwise, any description herein describing a determination of a value above a threshold and/or exceeding a predetermined or expected value is understood to include determination of a value below a threshold and/or falling below a predetermined or expected value.

[0619] The predetermined acceptable range may be based on anticipated system response or vibration based on the equipment geometry and control scheme such as number of bearings, relative rotational speed, influx of power to the system at a certain frequency, and the like. The predetermined acceptable range may also be based on long term analysis of detection values across a plurality of similar equipment and components and correlation of data with equipment failure. Based on vibration phase information, a physical location of a problem may be identified. Based on the vibration phase information system design flaws, off-nominal operation, and/or component or process failures may be identified. In example embodiments, an alert may be issued based on changes or rates of change in the data over time such as increasing amplitude or shifts in the

frequencies or phases at which a vibration occurs. In example embodiments, an alert may be issued based on accumulated values such as time spent over a threshold, weighted time spent over one or more thresholds, and/or an area of a curve of the detected value over one or more thresholds. In example embodiments, an alert may be issued based on a combination of data from different sensors such as relative changes in value, or relative rates of change in amplitude, frequency of phase in addition to values of non-phase sensors such as temperature, humidity and the like. For example, an increase in temperature and energy at certain frequencies may indicate a hot bearing that is starting to fail. In example embodiments, the relative criteria for an alarm may change with other data or information such as process stage, type of product being processed on equipment, ambient temperature and humidity, external vibrations from other equipment and the like.

[0620] In example embodiments, response circuit 8510 may cause the data acquisition circuit 8504 to enable or disable the processing of detection values corresponding to certain sensors based on the some of the criteria discussed above. This may include switching to sensors having different response rates, sensitivity, ranges, and the like; accessing new sensors or types of sensors, and the like. Switching may be undertaken based on a model, a set of rules, or the like. In example embodiments, switching may be under control of a machine learning system, such that switching is controlled based on one or more metrics of success, combined with input data, over a set of trials, which may occur under supervision of a human supervisor or under control of an automated system. Switching may involve switching from one input port to another (such as to switch from one sensor to another). Switching may involve altering the multiplexing of data, such as combining different streams under different circumstances. Switching may involve activating a system to obtain additional data, such as moving a mobile system (such as a robotic or drone system), to a location where different or additional data is available (such as positioning an image sensor for a different view or positioning a sonar sensor for a different direction of collection) or to a location where different sensors can be accessed (such as moving a collector to connect up to a sensor that is disposed at a location in an environment by a wired or wireless connection). The response circuit 8510 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 8510 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0621] In example embodiments, the response circuit 8510 may recommend maintenance at an upcoming process stop or initiate a maintenance call. The response circuit 8510 may recommend changes in process or operating parameters to remotely balance the piece of equipment. In example embodiments, the response circuit 8510 may implement or recommend process changes – for example to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but still at least partially operational, to change the operating speed of a component (such as to put it in a lower-demand mode), to initiate amelioration of an issue (such as to signal for additional lubrication of a roller bearing set, or to signal for an alignment process for a system that is out of balance), and the like.

[0622] In example embodiments, as shown in Figure 54, the data monitoring device 8540 may further comprise a data storage circuit 8542, memory, and the like. The signal evaluation circuit 8544 may periodically store certain detection values to enable the tracking of component performance over time.

[0623] In example embodiments, based on relevant operating conditions and/or failure modes which may occur in as sensor values approach one or more criteria, the signal evaluation circuit 8544 may store data in the data storage circuit 8542 based on the fit of data relative to one or more criteria, such as those described throughout this disclosure. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 8544 may store additional data such as RPMs, component loads, temperatures, pressures, vibrations or other sensor data of the types described throughout this disclosure. The signal evaluation circuit 8544 may store data at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0624] In example embodiments, as shown in Figure 55, a data monitoring system 8546 may comprise at least one data monitoring device 8548. The at least one data monitoring device 8548 comprising sensors 8506, a controller 8550 comprising a data acquisition circuit 8504, a signal evaluation circuit 8538, a data storage circuit 8542, and a communications circuit 8552 to allow data and analysis to be transmitted to a monitoring application 8556 on a remote server 8554. The signal evaluation circuit 8538 may comprise at least one of a phase detection circuit 8528, a phase lock loop circuit 8530, and/or a band pass circuit 8532. The signal evaluation circuit 8538 may periodically share data with the communication circuit 8552 for transmittal to the remote server 8554 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 8556. Because relevant operating conditions and/or failure modes may occur as sensor values approach one or more criteria, the signal evaluation circuit 8538 may share data with the communication circuit 8552 for transmittal to the remote server 8554 based on the fit of data relative to one or more criteria. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 8538 may share additional data such as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The signal evaluation circuit 8538 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server.

[0625] In example embodiments, as illustrated in Figure 56, a data collection system 8560 may have a plurality of monitoring devices 8558 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment (both the same and different types of equipment) in the same facility, as well as collecting data from monitoring devices in multiple facilities. A monitoring application on a remote server may receive and store the data coming from a plurality of the various monitoring devices. The monitoring application may then select subsets of data which may be jointly analyzed. Subsets of monitoring data may be selected based on data from a single type of component or data from a single type of equipment in which the component is operating. Monitoring data may be selected or

grouped based on common operating conditions such as size of load, operational condition (e.g., intermittent, continuous), operating speed or tachometer, common ambient environmental conditions such as humidity, temperature, air or fluid particulate, and the like. Monitoring data may be selected based on the effects of other nearby equipment, such as nearby machines rotating at similar frequencies, nearby equipment producing electromagnetic fields, nearby equipment producing heat, nearby equipment inducing movement or vibration, nearby equipment emitting vapors, chemicals or particulates, or other potentially interfering or intervening effects.

[0626] The monitoring application may then analyze the selected data set. For example, data from a single component may be analyzed over different time periods such as one operating cycle, several operating cycles, a month, a year, or the like. Data from multiple components of the same type may also be analyzed over different time periods. Trends in the data such as changes in frequency or amplitude may be correlated with failure and maintenance records associated with the same component or piece of equipment. Trends in the data such as changing rates of change associated with start-up or different points in the process may be identified. Additional data may be introduced into the analysis such as output product quality, output quantity (such as per unit of time), indicated success or failure of a process, and the like. Correlation of trends and values for different types of data may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected performance. This information may be transmitted back to the monitoring device to update types of data collected and analyzed locally or to influence the design of future monitoring devices.

[0627] In an illustrative and non-limiting example, the monitoring device may be used to collect and process sensor data to measure mechanical torque. The monitoring device may be in communication with or include a high resolution, high speed vibration sensor to collect data over an extended period of time, enough to measure multiple cycles of rotation. For gear driven equipment, the sampling resolution should be such that the number of samples taken per cycle is at least equal to the number of gear teeth driving the component. It will be understood that a lower sampling resolution may also be utilized, which may result in a lower confidence determination and/or taking data over a longer period of time to develop sufficient statistical confidence. This data may then be used in the generation of a phase reference (relative probe) or tachometer signal for a piece of equipment. This phase reference may be used to align phase data such as vibrational data or acceleration data from multiple sensors located at different positions on a component or on different components within a system. This information may facilitate the determination of torque for different components or the generation of an Operational Deflection Shape (ODS), indicating the extent of mechanical deflection of one or more components during an operational mode, which in turn may be used to measure mechanical torque in the component.

[0628] The higher resolution data stream may provide additional data for the detection of transitory signals in low speed operations. The identification of transitory signals may enable the identification of defects in a piece of equipment or component.

[0629] In an illustrative and non-limiting example, the monitoring device may be used to identify mechanical jitter for use in failure prediction models. The monitoring device may begin acquiring

data when the piece of equipment starts up through ramping up to operating speed and then during operation. Once at operating speed, it is anticipated that the torsional jitter should be minimal and changes in torsion during this phase may be indicative of cracks, bearing faults and the like. Additionally, known torsions may be removed from the signal to facilitate in the identification of unanticipated torsions resulting from system design flaws or component wear. Having phase information associated with the data collected at operating speed may facilitate identification of a location of vibration and potential component wear. Relative phase information for a plurality of sensors located throughout a machine may facilitate the evaluation of torsion as it is propagated through a piece of equipment.

[0630] An example system data collection in an industrial environment includes a data acquisition circuit that interprets a number of detection values from a number of input sensors communicatively coupled to the data acquisition circuit, each of the number of detection values corresponding to at least one of the input sensors, a signal evaluation circuit that obtains at least one of a vibration amplitude, a vibration frequency and a vibration phase location corresponding to at least one of the input sensors in response to the number of detection values, and a response circuit that performs at least one operation in response to at the at least one of the vibration amplitude, the vibration frequency and the vibration phase location. Certain further embodiments of an example system include: where the signal evaluation circuit includes a phase detection circuit, or a phase detection circuit and a phase lock loop circuit and/or a band pass filter; where the number of input sensors includes at least two input sensors providing phase information and at least one input sensor providing non-phase sensor information; the signal evaluation circuit further aligning the phase information provided by the at least two of the input sensors; where the at least one operation is further in response to at least one of: a change in magnitude of the vibration amplitude; a change in frequency or phase of vibration; a rate of change in at least one of vibration amplitude, vibration frequency and vibration phase; a relative change in value between at least two of vibration amplitude, vibration frequency and vibration phase; and/or a relative rate of change between at least two of vibration amplitude, vibration frequency, and vibration phase; the system further including an alert circuit, where the at least one operation includes providing an alert and where the alert may be one of haptic, audible and visual; a data storage circuit, where at least one of the vibration amplitude, vibration frequency, and vibration phase is stored periodically to create a vibration history, and where the at least one operation includes storing additional data in the data storage circuit (e.g., as a vibration fingerprint for a component); where the storing additional data in the data storage circuit is further in response to at least one of: a change in magnitude of the vibration amplitude; a change in frequency or phase of vibration; a rate of change in the vibration amplitude, frequency or phase; a relative change in value between at least two of vibration amplitude, frequency and phase; and a relative rate of change between at least two of vibration amplitude, frequency and phase; the system further comprising at least one of a multiplexing (MUX) circuit whereby alternative combinations of detection values may be selected based on at least one of user input, a detected state, and a selected operating parameter for a machine; where each of the number of detection values corresponds to at least one of the input sensors; where the

at least one operation includes enabling or disabling the connection of one or more portions of the multiplexing circuit; a MUX control circuit that interprets a subset of the number of detection values and provides the logical control of the MUX and the correspondence of MUX input and detected values as a result; and/or where the logic control of the MUX includes adaptive scheduling of the select lines.

[0631] An example method of monitoring a component, includes receiving time-based data from at least one sensor, phase-locking the received data with a reference signal, transforming the received time-based data to frequency data, filtering the frequency data to remove tachometer frequencies, identifying low amplitude signals occurring at high frequencies, and activating an alarm if a low amplitude signal exceeds a threshold.

[0632] An example system for data collection, processing, and utilization of signals in an industrial environment includes a plurality of monitoring devices, each monitoring device comprising a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and a vibration phase location corresponding to at least one of the input sensors in response to the corresponding at least one of the plurality of detection values; a data storage facility for storing a subset of the plurality of detection values; a communication circuit structured to communicate at least one selected detection value to a remote server; and a monitoring application on the remote server structured to: receive the at least one selected detection value; jointly analyze a subset of the detection values received from the plurality of monitoring devices; and recommend an action.

[0633] In certain further embodiments, an example system includes: for each monitoring device, the plurality of input sensors include at least one input sensor providing phase information and at least one input sensor providing non-phase input sensor information and where joint analysis includes using the phase information from the plurality of monitoring devices to align the information from the plurality of monitoring devices; where the subset of detection values is selected based on data associated with a detection value including at least one: common type of component, common type of equipment, and common operating conditions and further selected based on one of anticipated life of a component associated with detection values, type of the equipment associated with detection values, and operational conditions under which detection values were measured; and/or where the analysis of the subset of detection values includes feeding a neural net with the subset of detection values and supplemental information to learn to recognize various operating states, health states, life expectancies and fault states utilizing deep learning techniques, wherein the supplemental information comprises one of component specification, component performance, equipment specification, equipment performance, maintenance records, repair records and an anticipated state model.

[0634] An example system for data collection in an industrial environment includes a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection

values corresponding to at least one of the input sensors; a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to at least one of the input sensors in response to the corresponding at least one of a plurality of detection values; a multiplexing circuit whereby alternative combinations of the detection values may be selected based on at least one of user input, a detected state and a selected operating parameter for a machine, each of the plurality of detection values corresponding to at least one of the input sensors; and a response circuit structured to perform at least one operation in response to at the at least one of the vibration amplitude, vibration frequency and vibration phase location.

[0635] An example system for data collection in a piece of equipment, includes a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a response circuit structured to perform at least one operation in response to at the at least one of the vibration amplitude, vibration frequency and vibration phase location.

[0636] An example system for bearing analysis in an industrial environment, includes a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a life prediction comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value: and a response circuit structured to perform at least one operation in response to at the at least one of the vibration amplitude, vibration frequency and vibration phase location.

[0637] An example motor monitoring system includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the motor and motor components, store historical motor performance and buffer the plurality of detection values for a predetermined length of time; a timer

circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a motor analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a motor performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a motor performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and motor performance parameter.

[0638] An example system for estimating a vehicle steering system performance parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the vehicle steering system, the rack, the pinion, and the steering column, store historical steering system performance and buffer the plurality of detection values for a predetermined length of time;

[0639] a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a steering system analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a steering system performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a steering system performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the steering system performance parameter.

[0640] An example system for estimating a health parameter a pump performance parameter includes a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the pump and pump components associated with the detection values, store historical pump performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a pump analysis circuit structured to analyze buffered detection values relative to specifications and

anticipated state information resulting in a pump performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a pump performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the pump performance parameter, wherein the pump is one of a water pump in a car and a mineral pump.

[0641] An example system for estimating a drill performance parameter for a drilling machine, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the drill and drill components associated with the detection values, store historical drill performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a drill analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a drill performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a drill performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the drill performance parameter, wherein the drilling machine is one of an oil drilling machine and a gas drilling machine.

[0642] An example system for estimating a conveyor health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a conveyor and conveyor components associated with the detection values, store historical conveyor performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a conveyor analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a conveyor performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second

detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a conveyor performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the conveyor performance parameter.

[0643] An example system for estimating an agitator health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for an agitator and agitator components associated with the detection values, store historical agitator performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; an agitator analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in an agitator performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in an agitator performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the agitator performance parameter, wherein the agitator is one of a rotating tank mixer, a large tank mixer, a portable tank mixers, a tote tank mixer, a drum mixer, a mounted mixer and a propeller mixer.

[0644] An example system for estimating a compressor health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a compressor and compressor components associated with the detection values, store historical compressor performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a compressor analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a compressor performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a

signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a compressor performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the compressor performance parameter.

[0645] An example system for estimating an air conditioner health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for an air conditioner and air conditioner components associated with the detection values, store historical air conditioner performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; an air conditioner analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in an air conditioner performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in an air conditioner performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the air conditioner performance parameter.

[0646] An example system for estimating a centrifuge health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a centrifuge and centrifuge components associated with the detection values, store historical centrifuge performance and buffer the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a centrifuge analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a centrifuge performance parameter comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value and analyze the at least one of vibration

amplitude, vibration frequency and vibration phase location relative to buffered detection values, specifications and anticipated state information resulting in a centrifuge performance parameter; and a response circuit structured to perform at least one operation in response to at the at least one of vibration amplitude, vibration frequency and vibration phase location and the centrifuge performance parameter.

[0647] In example embodiments, information about the health of a component or piece of industrial equipment may be obtained by comparing the values of multiple signals at the same point in a process. This may be accomplished by aligning a signal relative to other related data signals, timers, or reference signals. An embodiment of a data monitoring device 8700, 8718 is shown in Figures 57-59 and may include a controller 8702, 8720. The controller may include a data acquisition circuit 8704, 8722, a signal evaluation circuit 8708, a data storage circuit 8716 and an optional response circuit 8710. The signal evaluation circuit 8708 may comprise a timer circuit 8714 and, optionally, a phase detection circuit 8712.

[0648] The data monitoring device may include a plurality of sensors 8706 communicatively coupled to a controller 8702. The plurality of sensors 8706 may be wired to ports on the data acquisition circuit 8704. The plurality of sensors 8706 may be wirelessly connected to the data acquisition circuit 8704 which may be able to access detection values corresponding to the output of at least one of the plurality of sensors 8706 where the sensors 8706 may be capturing data on different operational aspects of a piece of equipment or an operating component. In example embodiments, as illustrated in Figures 58 and 59, one or more external sensors 8724 which are not explicitly part of a monitoring device 8718 may be opportunistically connected to or accessed by the monitoring device 8718. The data acquisition circuit 8722 may include one or more input ports 8726. The one or more external sensors 8724 may be directly connected to the one or more input ports 8726 on the data acquisition circuit 8722 of the controller 8720. In example embodiments, as shown in Figure 59, a data acquisition circuit 8722 may further comprise a wireless communications circuit 8728 to access detection values corresponding to the one or more external sensors 8724 wirelessly or via a separate source or some combination of these methods.

[0649] The selection of the plurality of sensors 8706 8724 for connection to a data monitoring device 8700 8718 designed for a specific component or piece of equipment may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, resolution desired at various positions in a process or plant, reliability of the sensors, and the like. The impact of a failure, time response of a failure (e.g., warning time and/or off-nominal modes occurring before failure), likelihood of failure, and/or sensitivity required and/or difficulty to detect failed conditions may drive the extent to which a component or piece of equipment is monitored with more sensors and/or higher capability sensors being dedicated to systems where unexpected or undetected failure would be costly or have severe consequences.

[0650] The signal evaluation circuit 8708 may process the detection values to obtain information about a component or piece of equipment being monitored. Information extracted by the signal evaluation circuit 8708 may comprise information regarding what point or time in a process

corresponds with a detection value where the point in time is based on a timing signal generated by the timer circuit 8714. The start of the timing signal may be generated by detecting an edge of a control signal such as a rising edge, falling edge or both where the control signal may be associated with the start of a process. The start of the timing signal may be triggered by an initial movement of a component or piece of equipment. The start of the timing signal may be triggered by an initial flow through a pipe or opening or by a flow achieving a predetermined rate. The start of the timing signal may be triggered by a state value indicating a process has commenced – for example the state of a switch, button, data value provided to indicate the process has commenced, or the like. Information extracted may comprise information regarding a difference in phase, determined by the phase detection circuit 8712, between a stream of detection value and the time signal generated by the timer circuit 8714. Information extracted may comprise information regarding a difference in phase between one stream of detection values and a second stream of detection values where the first stream of detection values is used as a basis or trigger for a timing signal generated by the timer circuit.

[0651] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating and the like, sensors 8706 8724 may comprise one or more of, without limitation, a thermometer, a hygrometer, a voltage sensor, a current sensor, an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a displacement sensor, a turbidity meter, a viscosity meter, a load sensor, a tri-axial sensor, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an acoustical sensor, a pH sensor, and the like.

[0652] The sensors 8706 8724 may provide a stream of data over time that has a phase component, such as acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 8706 8724 may provide a stream of data that is not phase based such as temperature, humidity, load, and the like. The sensors 8706 8724 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0653] In example embodiments, as illustrated in Figures 60 and 61, the data acquisition circuit 8734 may further comprise a multiplexer circuit 8736 as described elsewhere herein. Outputs from the multiplexer circuit 8736 may be utilized by the signal evaluation circuit 8708. The response circuit 8710 may have the ability to turn on and off portions of the multiplexer circuit 8736. The response circuit 8710 may have the ability to control the control channels of the multiplexer circuit 8736.

[0654] The response circuit 8710 may further comprise evaluating the results of the signal evaluation circuit 8708 and, based on certain criteria, initiating an action. The criteria may include a sensor's detection values at certain frequencies or phases relative to the timer signal where the frequencies or phases of interest may be based on the equipment geometry, equipment control schemes, system input, historical data, current operating conditions, and/or an anticipated response.

Criteria may include a predetermined maximum or minimum value for a detection value from a specific sensor, a cumulative value of a sensor's corresponding detection value over time, a change in value, a rate of change in value, and/or an accumulated value (e.g., a time spent above/below a threshold value, a weighted time spent above/below one or more threshold values, and/or an area of the detected value above/below one or more threshold values). The criteria may comprise combinations of data from different sensors such as relative values, relative changes in value, relative rates of change in value, relative values over time, and the like. The relative criteria may change with other data or information such as process stage, type of product being processed, type of equipment, ambient temperature and humidity, external vibrations from other equipment, and the like.

[0655] Certain embodiments are described herein as detected values exceeding thresholds or predetermined values, but detected values may also fall below thresholds or predetermined values – for example where an amount of change in the detected value is expected to occur, but detected values indicate that the change may not have occurred. For example, and without limitation, vibrational data may indicate system agitation levels, properly operating equipment, or the like, and vibrational data below amplitude and/or frequency thresholds may be an indication of a process that is not operating according to expectations. Except where the context clearly indicates otherwise, any description herein describing a determination of a value above a threshold and/or exceeding a predetermined or expected value is understood to include determination of a value below a threshold and/or falling below a predetermined or expected value.

[0656] The predetermined acceptable range may be based on anticipated system response or vibration based on the equipment geometry and control scheme such as number of bearings, relative rotational speed, influx of power to the system at a certain frequency, and the like. The predetermined acceptable range may also be based on long term analysis of detection values across a plurality of similar equipment and components and correlation of data with equipment failure.

[0657] In example embodiments, an alert may be issued based on the some of the criteria discussed above. In an illustrative example, an increase in temperature and energy at certain frequencies may indicate a hot bearing that is starting to fail. In example embodiments, the relative criteria for an alarm may change with other data or information such as process stage, type of product being processed on equipment, ambient temperature and humidity, external vibrations from other equipment and the like. In an illustrative and non-limiting example, the response circuit 8710 may initiate an alert if a vibrational amplitude and/or frequency exceeds a predetermined maximum value, if there is a change or rate of change that exceeds a predetermined acceptable range, and/or if an accumulated value based on vibrational amplitude and/or frequency exceeds a threshold.

[0658] In example embodiments, response circuit 8710 may cause the data acquisition circuit 8704 to enable or disable the processing of detection values corresponding to certain sensors based on the some of the criteria discussed above. This may include switching to sensors having different response rates, sensitivity, ranges, and the like; accessing new sensors or types of sensors, and the like. This switching may be implemented by changing the control signals for a multiplexer circuit 8736 and/or by turning on or off certain input sections of the multiplexer circuit 8736. The response

circuit 8710 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 8710 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0659] In example embodiments, the response circuit 8710 may recommend maintenance at an upcoming process stop or initiate a maintenance call. The response circuit 8710 may recommend changes in process or operating parameters to remotely balance the piece of equipment. In example embodiments, the response circuit 8710 may implement or recommend process changes – for example to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but still at least partially operational. In an illustrative example, vibration phase information, derived by the phase detection circuit 8712 relative to a timer signal from the timer circuit 8714, may be indicative of a physical location of a problem. Based on the vibration phase information, system design flaws, off-nominal operation, and/or component or process failures may be identified.

[0660] In example embodiments, based on relevant operating conditions and/or failure modes which may occur in as sensor values approach one or more criteria, the signal evaluation circuit 8708 may store data in the data storage circuit 8716 based on the fit of data relative to one or more criteria. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 8708 may store additional data such as RPMs, component loads, temperatures, pressures, vibrations in the data storage circuit 8716. The signal evaluation circuit 8708 may store data at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0661] In example embodiments, as shown in Figures 62 and 63 and 64 and 65, a data monitoring system 8762 may include at least one data monitoring device 8768. The at least one data monitoring device 8768 may include sensors 8706 and a controller 8770 comprising a data acquisition circuit 8704, a signal evaluation circuit 8772, a data storage circuit 8742, and a communications circuit 8752 to allow data and analysis to be transmitted to a monitoring application 8776 on a remote server 8774. The signal evaluation circuit 8772 may include at least one of a phase detection circuit 8712 and a timer circuit 8714. The signal evaluation circuit 8772 may periodically share data with the communication circuit 8752 for transmittal to the remote server 8774 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 8776. Because relevant operating conditions and/or failure modes may occur as sensor values approach one or more criteria, the signal evaluation circuit 8708 may share data with the communication circuit 8752 for transmittal to the remote server 8774 based on the fit of data relative to one or more criteria. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 8708 may share additional data such as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The signal evaluation circuit 8772 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server.

[0662] In example embodiments, as shown in Figure 62, the communications circuit 8752 may communicate data directly to a remote server 8774. In example embodiments, as shown in Figure 63, the communications circuit 8752 may communicate data to an intermediate computer 8754 which may include a processor 8756 running an operating system 8758 and a data storage circuit 8760. The intermediate computer 8754 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 8774.

[0663] In embodiments as illustrated in Figures 64 and 65, a data collection system 8762 may have a plurality of monitoring devices 8768 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment, (both the same and different types of equipment) in the same facility as well as collecting data from monitoring devices in multiple facilities. In example embodiments, as show in in Figure 64 the communications circuit 8752 may communicated data directly to a remote server 8774. In example embodiments, as shown in Figure 65, the communications circuit 8752 may communicate data to an intermediate computer 8754 which may include a processor 8756 running an operating system 8758 and a data storage circuit 8760. The intermediate computer 8754 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 8774.

[0664] In example embodiments, a monitoring application 8776 on a remote server 8774 may receive and store one or more of detection values, timing signals and data coming from a plurality of the various monitoring devices 8768. The monitoring application 8776 may then select subsets of the detection values, timing signals and data to be jointly analyzed. Subsets for analysis may be selected based on a single type of component or a single type of equipment in which a component is operating. Subsets for analysis may be selected or grouped based on common operating conditions such as size of load, operational condition (e.g., intermittent, continuous, process stage), operating speed or tachometer, common ambient environmental conditions such as humidity, temperature, air or fluid particulate, and the like. Subsets for analysis may be selected based on the effects of other nearby equipment such as nearby machines rotating at similar frequencies.

[0665] The monitoring application 8776 may then analyze the selected subset. In an illustrative example, data from a single component may be analyzed over different time periods such as one operating cycle, several operating cycles, a month, a year, the life of the component or the like. Data from multiple components of the same type may also be analyzed over different time periods. Trends in the data such as changes in frequency or amplitude may be correlated with failure and maintenance records associated with the same or a related component or piece of equipment. Trends in the data such as changing rates of change associated with start-up or different points in the process may be identified. Additional data may be introduced into the analysis such as output product quality, indicated success or failure of a process, and the like. Correlation of trends and values for different types of data may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected performance. This information may be transmitted back to the monitoring device to update types of data collected and analyzed locally or to influence the design of future monitoring devices.

[0666] In an illustrative and non-limiting example, a monitoring device 8768 may be used to collect and process sensor data to measure mechanical torque. The monitoring device 8768 may be in communication with or include a high resolution, high speed vibration sensor to collect data over a period of time sufficient to measure multiple cycles of rotation. For gear driven components, the sampling resolution of the sensor should be such that the number of samples taken per cycle is at least equal to the number of gear teeth driving the component. It will be understood that a lower sampling resolution may also be utilized, which may result in a lower confidence determination and/or taking data over a longer period of time to develop sufficient statistical confidence. This data may then be used in the generation of a phase reference (relative probe) or tachometer signal for a piece of equipment. This phase reference may be used directly or used by the timer circuit 8714 to generate a timing signal to align phase data such as vibrational data or acceleration data from multiple sensors located at different positions on a component or on different components within a system. This information may facilitate the determination of torque for different components or the generation of an Operational Deflection Shape (ODS).

[0667] A higher resolution data stream may also provide additional data for the detection of transitory signals in low speed operations. The identification of transitory signals may enable the identification of defects in a piece of equipment or component operating at low RPMs.

[0668] In an illustrative and non-limiting example, the monitoring device may be used to identify mechanical jitter for use in failure prediction models. The monitoring device may begin acquiring data when the piece of equipment starts up, through ramping up to operating speed, and then during operation. Once at operating speed, it is anticipated that the torsional jitter should be minimal or within expected ranges, and changes in torsion during this phase may be indicative of cracks, bearing faults, and the like. Additionally, known torsions may be removed from the signal to facilitate the identification of unanticipated torsions resulting from system design flaws, component wear, or unexpected process events. Having phase information associated with the data collected at operating speed may facilitate identification of a location of vibration and potential component wear, and/or may be further correlated to a type of failure for a component. Relative phase information for a plurality of sensors located throughout a machine may facilitate the evaluation of torsion as it is propagated through a piece of equipment.

[0669] In example embodiments, the monitoring application 8776 may have access to equipment specifications, equipment geometry, component specifications, component materials, anticipated state information for plurality of component types, operational history, historical detection values, component life models, and the like for use in analyzing the selected subset using rule-based or model-based analysis. In example embodiments, the monitoring application 8776 may feed a neural net with the selected subset to learn to recognize various operating state, health states (e.g., lifetime predictions) and fault states utilizing deep learning techniques. In example embodiments, a hybrid of the two techniques (model-based learning and deep learning) may be used.

[0670] In an illustrative and non-limiting example, component health of: conveyors and lifters in an assembly line; water pumps on industrial vehicles; factory air conditioning units; drilling machines, screw drivers, compressors, pumps, gearboxes, vibrating conveyors, mixers and motors

situated in the oil and gas fields; factory mineral pumps; centrifuges, and refining tanks situated in oil and gas refineries; and compressors in gas handling systems may be monitored using the phase detection and alignment techniques, data monitoring devices and data collection systems described herein.

[0671] In an illustrative and non-limiting example, the component health of equipment to promote chemical reactions deployed in chemical and pharmaceutical production lines (e.g., rotating tank/mixer agitators, mechanical/rotating agitators, and propeller agitators) may be evaluated using the phase detection and alignment techniques, data monitoring devices and data collection systems described herein.

[0672] In an illustrative and non-limiting example, the component health of vehicle steering mechanisms and/or vehicle engines may be evaluated using the phase detection and alignment techniques, data monitoring devices and data collection systems described herein.

[0673] An example monitoring system for data collection, includes a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a signal evaluation circuit comprising: a timer circuit structured to generate at least one timing signal; and a phase detection circuit structured to determine a relative phase difference between at least one of the plurality of detection values and at least one of the timing signals from the timer circuit; and a response circuit structured to perform at least one operation in response to the relative phase difference. In certain further embodiments, an example system includes:

[0674] wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; a data storage circuit, wherein the relative phase difference and at least one of the detection values and the timing signal are stored; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference; wherein the data acquisition circuit further comprises at least one multiplexer circuit (MUX) whereby alternative combinations of detection values may be selected based on at least one of user input and a selected operating parameter for a machine, wherein each of the plurality of detection values corresponds to at least one of the input sensors; wherein the at least one operation comprises enabling or disabling one or more portions of the multiplexer circuit, or altering the multiplexer control lines; wherein the data acquisition circuit comprises at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits; and/or the system further comprising a MUX control circuit structured to interpret a subset of the plurality of detection values and provide

the logical control of the MUX and the correspondence of MUX input and detected values as a result, wherein the logic control of the MUX comprises adaptive scheduling of the select lines.

[0675] An example system for data collection, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a signal evaluation circuit comprising: a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; and a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a phase response circuit structured to perform at least one operation in response to the phase difference. In certain further embodiments, an example system includes wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one the plurality of detection values and a relative rate of change in amplitude and relative phase of at least one the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; where the system, further includes a data storage circuit; wherein the relative phase difference and at least one of the detection values and the timing signal are stored; wherein the at least one operation further includes storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference; wherein the data acquisition circuit further includes at least one multiplexer (MUX) circuit whereby alternative combinations of detection values may be selected based on at least one of user input and a selected operating parameter for a machine; wherein each of the plurality of detection values corresponds to at least one of the input sensors; wherein the at least one operation comprises enabling or disabling one or more portions of the multiplexer circuit, or altering the multiplexer control lines; wherein the data acquisition circuit comprises at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits; where the system further comprising a MUX control circuit structured to interpret a subset of the plurality of detection values and provide the logical control of the MUX and the correspondence of MUX input and detected values as a result; and/or wherein the logic control of the MUX comprises adaptive scheduling of the select lines.

[0676] An example system for data collection, processing, and utilization of signals in an industrial environment includes a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a signal evaluation circuit comprising: a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; and a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; a data storage facility for storing a subset of the plurality of detection values and the

timing signal; a communication circuit structured to communicate at least one selected detection value and the timing signal to a remote server; and a monitoring application on the remote server structured to receive the at least one selected detection value and the timing signal; jointly analyze a subset of the detection values received from the plurality of monitoring devices; and recommend an action. In certain example embodiments, the example system further includes wherein joint analysis comprises using the timing signal from each of the plurality of monitoring devices to align the detection values from the plurality of monitoring devices and/or wherein the subset of detection values is selected based on data associated with a detection value comprising at least one: common type of component, common type of equipment, and common operating conditions.

[0677] An example system for data collection in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit, the data acquisition circuit comprising a multiplexer circuit whereby alternative combinations of the detection values may be selected based on at least one of user input, a detected state and a selected operating parameter for a machine, each of the plurality of detection values corresponding to at least one of the input sensors; a signal evaluation circuit comprising: a timer circuit structured to generate a timing signal; and a phase detection circuit structured to determine a relative phase difference between at least one of the plurality of detection values and a signal from the timer circuit; and a response circuit structured to perform at least one operation in response to the phase difference.

[0678] An example monitoring system for data collection in a piece of equipment, includes a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; and a response circuit structured to perform at least one operation in response to at the at least one of the vibration amplitude, vibration frequency and vibration phase location.

[0679] A monitoring system for bearing analysis in an industrial environment, the monitoring device includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a timer circuit structured to generate a timing signal a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; a timer circuit structured to generate a timing signal based on a first detected value of the plurality of detection values; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a life prediction

comprising: a phase detection circuit structured to determine a relative phase difference between a second detection value of the plurality of detection values and the timing signal; a signal evaluation circuit structured to obtain at least one of vibration amplitude, vibration frequency and vibration phase location corresponding to a second detected value: and a response circuit structured to perform at least one operation in response to at the at least one of the vibration amplitude, vibration frequency and vibration phase location.

[0680] In example embodiments, information about the health or other status or state information of or regarding a component or piece of industrial equipment may be obtained by monitoring the condition of various components throughout a process. Monitoring may include monitoring the amplitude of a sensor signal measuring attributes such as temperature, humidity, acceleration, displacement and the like. An embodiment of a data monitoring device 9000 is shown in Figure 66 and may include a plurality of sensors 9006 communicatively coupled to a controller 9002. The controller 9002, which may be part of a data collection device, such as a mobile data collector, or part of a system, such as a network-deployed or cloud-deployed system, may include a data acquisition circuit 9004, a signal evaluation circuit 9008 and a response circuit 9010. The signal evaluation circuit 9008 may comprise a peak detection circuit 9012. Additionally, the signal evaluation circuit 9008 may optionally comprise one or more of a phase detection circuit 9016, a bandpass filter circuit 9018, a phase lock loop circuit, a torsional analysis circuit, a bearing analysis circuit, and the like. The bandpass filter 9018 may be used to filter a stream of detection values such that values, such as peaks and valleys, are detected only at or within bands of interest, such as frequencies of interest. The data acquisition circuit 9004 may include one or more analog-to-digital converter circuits 9014. A peak amplitude detected by the peak detection circuit 9012 may be input into one or more analog-to-digital converter circuits 9014 to provide a reference value for scaling output of the analog-to-digital converter circuits 9014 appropriately.

[0681] The plurality of sensors 9006 may be wired to ports on the data acquisition circuit 9004. The plurality of sensors 9006 may be wirelessly connected to the data acquisition circuit 9004. The data acquisition circuit 9004 may be able to access detection values corresponding to the output of at least one of the plurality of sensors 9006 where the sensors 9006 may be capturing data on different operational aspects of a piece of equipment or an operating component.

[0682] The selection of the plurality of sensors 9006 for a data monitoring device 9000 designed for a specific component or piece of equipment may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, resolution desired at various positions in a process or plant, reliability of the sensors, power availability, power utilization, storage utilization, and the like. The impact of a failure, time response of a failure (e.g., warning time and/or off-optimal modes occurring before failure), likelihood of failure, extent of impact of failure, and/or sensitivity required and/or difficulty to detection failure conditions may drive the extent to which a component or piece of equipment is monitored with more sensors and/or higher capability sensors being dedicated to systems where unexpected or undetected failure would be costly or have severe consequences.

[0683] The signal evaluation circuit 9008 may process the detection values to obtain information about a component or piece of equipment being monitored. Information extracted by the signal evaluation circuit 9008 may comprise information regarding a peak value of a signal such as a peak temperature, peak acceleration, peak velocity, peak pressure, peak weight bearing, peak strain, peak bending, or peak displacement. The peak detection may be done using analog or digital circuits. In example embodiments, the peak detection circuit 9012 may be able to distinguish between "local" or short term peaks in a stream of detection values and a "global" or longer term peak. In example embodiments, the peak detection circuit 9012 may be able to identify peak shapes (not just a single peak value) such as flat tops, asymptotic approaches, discrete jumps in the peak value or rapid/steep climbs in peak value, sinusoidal behavior within ranges and the like. Flat topped peaks may indicate saturation of a sensor. Asymptotic approaches to a peak may indicate linear system behavior. Discrete jumps in value or steep changes in peak value may indicate quantized or nonlinear behavior of either the sensor doing the measurement or the behavior of the component. In example embodiments, the system may be able to identify sinusoidal variations in the peak value within an envelope, such as an envelope established by line or curve connecting a series of peak values. It should be noted that references to "peaks" should be understood to encompass one or more "valleys," representing a series of low points in measurement, except where context indicates otherwise.

[0684] In example embodiments, a peak value may be used as a reference for an analog-to-digital conversion circuit 9014.

[0685] In an illustrative and non-limiting example, a temperature probe may measure the temperature of a gear as it rotates in a machine. The peak temperature may be detected by a peak detection circuit 9012. The peak temperature may be fed into an analog-to-digital converter circuit 9014 to appropriately scale a stream of detection values corresponding to temperature readings of the gear as it rotates in a machine. The phase of the stream of detection values corresponding to temperature relative to an orientation of the gear may be determined by the phase detection circuit 9016. Knowing where in the rotation of the gear a peak temperature is occurring may allow the identification of a bad gear tooth.

[0686] In example embodiments, two or more sets of detection values may be fused to create detection values for a virtual sensor. A peak detection circuit may be used to verify consistency in timing of peak values between at least one of the two or more sets of detection values and the detection values for the virtual sensor.

[0687] In example embodiments, the signal evaluation circuit 9008 may be able to reset the peak detection circuit 9012 upon start-up of the monitoring device 9000, upon edge detection of a control signal of the system being monitored, based on a user input, after a system error and the like. In example embodiments, the signal evaluation circuit 9008 may discard an initial portion of the output of the peak detection circuit 9012 prior to using the peak value as a reference value for an analog-to-digital conversion circuit to allow the system to fully come on line.

[0688] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating and the like, sensors 9006 may comprise one or more of, without

limitation, a vibration sensor, a thermometer, a hygrometer, a voltage sensor, a current sensor, an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, an acoustic wave sensor, a displacement sensor, a turbidity meter, a viscosity meter, a load sensor, a tri-axial sensor, an accelerometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an acoustical sensor, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference.

[0689] The sensors 9006 may provide a stream of data over time that has a phase component, such as relating to acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 9006 may provide a stream of data that is not conventionally phase-based, such as temperature, humidity, load, and the like. The sensors 9006 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0690] In example embodiments, as illustrated in Figure 66, the sensors 9006 may be part of the data monitoring device 9000, referred to herein in some cases as a data collector, which in some cases may comprise a mobile or portable data collector. In example embodiments, as illustrated in Figures 67 and 68, one or more external sensors 9026, which are not explicitly part of a monitoring device 9020 but rather are new, previously attached to or integrated into the equipment or component, may be opportunistically connected to or accessed by the monitoring device 9020. The monitoring device 9020 may include a controller 9022. The controller 9022 may include a response circuit 9010, a signal evaluation circuit 9008 and a data acquisition circuit 9024. The signal evaluation circuit 9008 may include a peak detection circuit 9012 and optionally a phase detection circuit 9016 and/or a bandpass filter circuit 9018. The data acquisition circuit 9024 may include one or more input ports 9028. The one or more external sensors 9026 may be directly connected to the one or more input ports 9028 on the data acquisition circuit 9024 of the controller 9022 or may be accessed by the data acquisition circuit 9004 wirelessly, such as by a reader, interrogator, or other wireless connection, such as over a short-distance wireless protocol. In embodiments as shown in Figure 68, a data acquisition circuit 9024 may further comprise a wireless communication circuit 9030. The data acquisition circuit 9024 may use the wireless communication circuit 9030 to access detection values corresponding to the one or more external sensors 9026 wirelessly or via a separate source or some combination of these methods.

[0691] In embodiments as illustrated in Figure 69, the data acquisition circuit 9036 may further comprise a multiplexer circuit 9038 as described elsewhere herein. Outputs from the multiplexer circuit 9038 may be utilized by the signal evaluation circuit 9008. The response circuit 9010 may have the ability to turn on and off portions of the multiplexor circuit 9038. The response circuit 9010 may have the ability to control the control channels of the multiplexor circuit 9038.

[0692] The response circuit 9010 may evaluate the results of the signal evaluation circuit 9008 and, based on certain criteria, initiate an action. The criteria may include a predetermined peak

value for a detection value from a specific sensor, a cumulative value of a sensor's corresponding detection value over time, a change in peak value, a rate of change in a peak value, and/or an accumulated value (e.g., a time spent above/below a threshold value, a weighted time spent above/below one or more threshold values, and/or an area of the detected value above/below one or more threshold values). The criteria may comprise combinations of data from different sensors such as relative values, relative changes in value, relative rates of change in value, relative values over time, and the like. The relative criteria may change with other data or information such as process stage, type of product being processed, type of equipment, ambient temperature and humidity, external vibrations from other equipment, and the like. The relative criteria may be reflected in one or more calculated statistics or metrics (including ones generated by further calculations on multiple criteria or statistics), which in turn may be used for processing (such as an on-board a data collector or by an external system), such as to be provided as an input to one or more of the machine learning capabilities described in this disclosure, to a control system (which may be on-board a data collector or remote, such as to control selection of data inputs, multiplexing of sensor data, storage, or the like), or as a data element that is an input to another system, such as a data stream or data package that may be available to a data marketplace, a SCADA system, a remote control system, a maintenance system, an analytic system, or other system.

[0693] Certain embodiments are described herein as detected values exceeding thresholds or predetermined values, but detected values may also fall below thresholds or predetermined values – for example where an amount of change in the detected value is expected to occur, but detected values indicate that the change may not have occurred. For example, and without limitation, vibrational data may indicate system agitation levels, properly operating equipment, or the like, and vibrational data below amplitude and/or frequency thresholds may be an indication of a process that is not operating according to expectations. For example, in a process involving a blender, a mixer, an agitator or the like, the absence of vibration may indicate that a blade, fin, vane or other working element is unable to move adequately, such as, for example, as a result of a working material being excessively viscous or as a result of a problem in gears (e.g., stripped gears, seizing in gears, or the like (a clutch, or the like). Except where the context clearly indicates otherwise, any description herein describing a determination of a value above a threshold and/or exceeding a predetermined or expected value is understood to include determination of a value below a threshold and/or falling below a predetermined or expected value.

[0694] The predetermined acceptable range may be based on anticipated system response or vibration based on the equipment geometry and control scheme such as number of bearings, relative rotational speed, influx of power to the system at a certain frequency, and the like. The predetermined acceptable range may also be based on long term analysis of detection values across a plurality of similar equipment and components and correlation of data with equipment failure.

[0695] In example embodiments, the response circuit 9010 may issue an alert based on one or more of the criteria discussed above. In an illustrative example, an increase in peak temperature beyond a predetermined value may indicate a hot bearing that is starting to fail. In example embodiments, the relative criteria for an alarm may change with other data or information such as

process stage, type of product being processed on equipment, ambient temperature and humidity, external vibrations from other equipment and the like. In an illustrative and non-limiting example, the response circuit 9010 may initiate an alert if an amplitude, such as a vibrational amplitude and/or frequency, exceeds a predetermined maximum value, if there is a change or rate of change that exceeds a predetermined acceptable range, and/or if an accumulated value based on such amplitude and/or frequency exceeds a threshold.

[0696] In example embodiments, the response circuit 9010 may cause the data acquisition circuit 9004 to enable or disable the processing of detection values corresponding to certain sensors based on one or more of the criteria discussed above. This may include switching to sensors having different response rates, sensitivity, ranges, and the like; accessing new sensors or types of sensors, accessing data from multiple sensors, and the like. Switching may be based on a detected peak value for the sensor being switched or based on the peak value of another sensor. Switching may be undertaken based on a model, a set of rules, or the like. In example embodiments, switching may be under control of a machine learning system, such that switching is controlled based on one or more metrics of success, combined with input data, over a set of trials, which may occur under supervision of a human supervisor or under control of an automated system. Switching may involve switching from one input port to another (such as to switch from one sensor to another). Switching may involve altering the multiplexing of data, such as combining different streams under different circumstances. Switching may involve activating a system to obtain additional data, such as moving a mobile system (such as a robotic or drone system), to a location where different or additional data is available (such as positioning an image sensor for a different view or positioning a sonar sensor for a different direction of collection) or to a location where different sensors can be accessed (such as moving a collector to connect up to a sensor that is disposed at a location in an environment by a wired or wireless connection). This switching may be implemented by changing the control signals for a multiplexor circuit 9038 and/or by turning on or off certain input sections of the multiplexor circuit 9038.

[0697] In example embodiments, the response circuit 9010 may adjust a sensor scaling value using the detected peak as a reference voltage. The response circuit 9010 may adjust a sensor sampling rate such that the peak value is captured.

[0698] The response circuit 9010 may identify sensor overload. In example embodiments, the response circuit 9010 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 9010 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0699] In example embodiments, the response circuit 9010 may recommend maintenance at an upcoming process stop or initiate a maintenance call where the maintenance may include the replacement of the sensor with the same or an alternate type of sensor having a different response rate, sensitivity, range and the like. In example embodiments, the response circuit 9010 may implement or recommend process changes— for example, to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but still at least

partially operational, to change the operating speed of a component (such as to put it in a lower-demand mode), to initiate amelioration of an issue (such as to signal for additional lubrication of a roller bearing set, or to signal for an alignment process for a system that is out of balance), and the like.

[0700] In example embodiments, as shown in Figure 70, the data monitoring device 9040 may include sensors 9006 and a controller 9042 which may include a data acquisition circuit 9004, and a signal evaluation circuit 9008. The signal evaluation circuit 9008 may include a peak detection circuit 9012 and, optionally, a phased detection circuit 9016 and/or a bandpass filter circuit 9018. The controller 9042 may further include a data storage circuit 9044, memory, and the like. The controller 9042 may further include a response circuit 9010. The signal evaluation circuit 9008 may periodically store certain detection values in the data storage circuit 9044 to enable the tracking of component performance over time.

[0701] In example embodiments, based on relevant criteria as described elsewhere herein, operating conditions and/or failure modes which may occur as sensor values approach one or more criteria, the signal evaluation circuit 9008 may store data in the data storage circuit 9044 based on the fit of data relative to one or more criteria, such as those described throughout this disclosure. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 9008 may store additional data such as RPMs, component loads, temperatures, pressures, vibrations or other sensor data of the types described throughout this disclosure in the data storage circuit 9068. The signal evaluation circuit 9008 may store data at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0702] In example embodiments, the signal evaluation circuit 9008 may store new peaks that indicate changes in overall scaling over a long duration (e.g., scaling a data stream based on historical peaks over months of analysis). The signal evaluation circuit 9008 may store data when historical peak values are approached (e.g., as temperatures, pressures, vibrations, velocities, accelerations and the like approach historical peaks).

[0703] In embodiments as shown in Figures 71 and 72 and 73 and 74, a data monitoring system 9046 may include at least one data monitoring device 9048. At least one data monitoring device 9048 may include sensors 9006 and a controller 9050 comprising a data acquisition circuit 9004, a signal evaluation circuit 9008, a data storage circuit 9044, and a communication circuit 9052 to allow data and analysis to be transmitted to a monitoring application 9056 on a remote server 9054. The signal evaluation circuit 9008 may include at least one of a peak detection circuit 9012. The signal evaluation circuit 9008 may periodically share data with the communication circuit 9052 for transmittal to the remote server 9054 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 9056. Because relevant operating conditions and/or failure modes may occur as sensor values approach one or more criteria as described elsewhere herein, the signal evaluation circuit 9008 may share data with the communication circuit 9052 for transmittal to the remote server 9054 based on the fit of data

relative to one or more criteria. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 9008 may share additional data such as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The signal evaluation circuit 9008 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server.

[0704] In example embodiments, as shown in Figure 71, the communication circuit 9052 may communicate data directly to a remote server 9054. In example embodiments, as shown in Figure 72, the communication circuit 9052 may communicate data to an intermediate computer 9058 which may include a processor 9060 running an operating system 9062 and a data storage circuit 9064.

[0705] In example embodiments, as illustrated in Figures 73 and 74, a data collection system 9066 may have a plurality of monitoring devices 9048 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment (both the same and different types of equipment) in the same facility as well as collecting data from monitoring devices in multiple facilities. A monitoring application 9056 on a remote server 9054 may receive and store one or more of detection values, timing signals or data coming from a plurality of the various monitoring devices 9048.

[0706] In example embodiments, as shown in Figure 71, the communication circuit 9052 may communicate data directly to a remote server 9054. In example embodiments, as shown in Figure 72, the communication circuit 9052 may communicate data to an intermediate computer 9058 which may include a processor 9060 running an operating system 9062 and a data storage circuit 9064. There may be an individual intermediate computer 9058 associated with each monitoring device 9048 or an individual intermediate computer 9058 may be associated with a plurality of monitoring devices 9048 where the intermediate computer 9058 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 9054.

[0707] The monitoring application 9056 may select subsets of the detection values, timing signals and data to be jointly analyzed. Subsets for analysis may be selected based on a single type of component or a single type of equipment in which a component is operating. Subsets for analysis may be selected or grouped based on common operating conditions such as size of load, operational condition (e.g., intermittent, continuous), operating speed or tachometer, common ambient environmental conditions such as humidity, temperature, air or fluid particulate, and the like. Subsets for analysis may be selected based on the effects of other nearby equipment such as nearby machines rotating at similar frequencies, nearby equipment producing electromagnetic fields, nearby equipment producing heat, nearby equipment inducing movement or vibration, nearby equipment emitting vapors, chemicals or particulates, or other potentially interfering or intervening effects.

[0708] The monitoring application 9056 may then analyze the selected subset. In an illustrative example, data from a single component may be analyzed over different time periods such as one operating cycle, several operating cycles, a month, a year, the life of the component or the like. Data from multiple components of the same type may also be analyzed over different time periods.

Trends in the data such as changes in frequency or amplitude may be correlated with failure and maintenance records associated with the same or a related component or piece of equipment. Trends in the data, such as changing rates of change associated with start-up or different points in the process, may be identified. Additional data may be introduced into the analysis such as output product quality, output quantity (such as per unit of time), indicated success or failure of a process, and the like. Correlation of trends and values for different types of data may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected performance. This information may be transmitted back to the monitoring device to update types of data collected and analyzed locally or to influence the design of future monitoring devices.

[0709] In example embodiments, the monitoring application 9056 may have access to equipment specifications, equipment geometry, component specifications, component materials, anticipated state information for a plurality of component types, operational history, historical detection values, component life models and the like for use analyzing the selected subset using rule-based or model-based analysis. In example embodiments, the monitoring application 9056 may feed a neural net with the selected subset to learn to recognize peaks in waveform patterns by feeding a large data set sample of waveform behavior of a given type within which peaks are designated (such as by human analysts).

[0710] A monitoring system for data collection in an industrial environment, the monitoring system comprising: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a peak detection circuit structured to determine at least one peak value in response to the plurality of detection values; and a peak response circuit structured to perform at least one operation in response to the at least one peak value.

[0711] An example monitoring system further includes: wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values' wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible or visual; further comprising a data storage circuit, wherein the relative phase difference and at least one of the detection values and the timing signal are stored wherein the at least one operation further comprises storing additional data in the data storage circuit wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference wherein the data acquisition circuit further comprises at least one multiplexer circuit whereby alternative combinations of detection values may be selected based on at least one of user input and a selected operating parameter for a machine, wherein each of the plurality of detection values corresponds to at least one of the input sensors wherein the at least one operation comprises enabling or disabling one or more portions of the multiplexer circuit, or altering the multiplexer

control lines wherein the data acquisition circuit comprises at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits.

[0712] A monitoring system for data collection in an industrial environment, the monitoring system structure to receive input corresponding to a plurality of sensors, includes a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input sensors; a peak detection circuit structured to determine at least one peak value in response to the plurality of detection values; and a peak response circuit structured to perform at least one operation in response to the at least one peak value.

[0713] An example monitoring system further includes: wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values wherein the at least one operation comprises issuing an alert wherein the alert may be one of haptic, audible or visual further comprising a data storage circuit, wherein the relative phase difference and at least one of the detection values and the timing signal are stored wherein the at least one operation further comprises storing additional data in the data storage circuit wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference wherein the data acquisition circuit further comprises at least one multiplexer circuit whereby alternative combinations of detection values may be selected based on at least one of user input and a selected operating parameter for a machine, wherein each of the plurality of detection values corresponds to at least one of the input sensors wherein the at least one operation comprises enabling or disabling one or more portions of the multiplexer circuit, or altering the multiplexer control lines wherein the data acquisition circuit comprises at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits.

[0714] An example system for data collection, processing, and utilization of signals in an industrial environment includes: a plurality of monitoring devices, each monitoring device comprising: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a peak detection circuit structured to determine at least one peak value in response to the plurality of detection values; a peak response circuit structured to select at least one detection value in response to the at least one peak value; a communication circuit structured to communicate the at least one selected detection value to a remote server; and a monitoring application on the remote server structured to: receive the at least one selected detection value; jointly analyze received detection values from a subset of the plurality of monitoring devices; and recommend an action.

[0715] An example system further includes: the system further structured to subset detection values based on one of anticipated life of a component associated with detection values, type of the equipment associated with detection values, and operational conditions under which detection values were measured; wherein the analysis of the subset of detection values comprises feeding a neural net with the subset of detection values and supplemental information to learn to recognize various operating states, health states, life expectancies and fault states utilizing deep learning techniques; wherein the supplemental information comprises one of component specification, component performance, equipment specification, equipment performance, maintenance records, repair records and an anticipated state model wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one the plurality of detection values wherein the at least one operation comprises issuing an alert wherein the alert may be one of haptic, audible and visual further comprising a data storage circuit, wherein the relative phase difference and at least one of the detection values and the timing signal are stored wherein the at least one operation further comprises storing additional data in the data storage circuit wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference wherein the data acquisition circuit further comprises at least one multiplexer circuit whereby alternative combinations of detection values may be selected based on at least one of user input and a selected operating parameter for a machine, wherein each of the plurality of detection values corresponds to at least one of the input sensors wherein the at least one operation comprises enabling or disabling one or more portions of the multiplexer circuit, or altering the multiplexer control lines and/or wherein the data acquisition circuit comprises at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits.

[0716] An example motor monitoring system, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the motor and motor components, store historical motor performance and buffer the plurality of detection values for a predetermined length of time;

a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a motor performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and a motor system performance parameter.

[0717] An example system for estimating a vehicle steering system performance parameter, the device includes: a data acquisition circuit structured to interpret a plurality of detection values from

a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the vehicle steering system, the rack, the pinion, and the steering column, store historical steering system performance and buffer the plurality of detection values for a predetermined length of time; a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a vehicle steering system performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and a vehicle steering system performance parameter.

[0718] An example system for estimating a pump performance parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the pump and pump components associated with the detection values, store historical pump performance and buffer the plurality of detection values for a predetermined length of time; a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a pump performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and a pump performance parameter. In certain further embodiments, the example system includes wherein the pump is a water pump in a car and wherein the pump is a mineral pump.

[0719] An example system for estimating a drill performance parameter for a drilling machine, includes a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the drill and drill components associated with the detection values, store historical drill performance and buffer the plurality of detection values for a predetermined length of time; a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a drill performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and a drill performance parameter. An example system further includes wherein the drilling machine is one of an oil drilling machine and a gas drilling machine.

[0720] An example system for estimating a conveyor health parameter, the system includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a conveyor and conveyor components associated with the detection values, store historical conveyor performance and buffer the plurality of detection values for a predetermined length of time; a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a conveyor performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and a conveyor performance parameter.

[0721] An example system for estimating an agitator health parameter, the system includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for an agitator and agitator components associated with the detection values, store historical agitator performance and buffer the plurality of detection values for a predetermined length of time; a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in an agitator performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and an agitator performance parameter. In certain example embodiments, a system further includes where the agitator is one of a rotating tank mixer, a large tank mixer, a portable tank mixer, a tote tank mixer, a drum mixer, a mounted mixer and a propeller mixer.

[0722] An example system for estimating a compressor health parameter, the system includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a compressor and compressor components associated with the detection values, store historical compressor performance and buffer the plurality of detection values for a predetermined length of time;

a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a compressor performance parameter; and a peak

response circuit structured to perform at least one operation in response to one of a peak value and a compressor performance parameter.

[0723] An example system for estimating an air conditioner health parameter, the system includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for an air conditioner and air conditioner components associated with the detection values, store historical air conditioner performance and buffer the plurality of detection values for a predetermined length of time; a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value, a pressure value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in an air conditioner performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and an air conditioner performance parameter.

[0724] An example system for estimating a centrifuge health parameter, the system includes: a data acquisition circuit structured to interpret a plurality of detection values from a plurality of input sensors communicatively coupled to the data acquisition circuit, each of the plurality of detection values corresponding to at least one of the input sensors; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a centrifuge and centrifuge components associated with the detection values, store historical centrifuge performance and buffer the plurality of detection values for a predetermined length of time;

a peak detection circuit structured to determine a plurality of peak values comprising at least a temperature peak value, a speed peak value and a vibration peak value in response to the plurality of detection values and analyze the peak values relative to buffered detection values, specifications and anticipated state information resulting in a centrifuge performance parameter; and a peak response circuit structured to perform at least one operation in response to one of a peak value and a centrifuge performance parameter.

[0725] Bearings are used throughout many different types of equipment and applications. Bearings may be present in or supporting shafts, motors, rotors, stators, housings, frames, suspension systems and components, gears, gear sets of various types, other bearings, and other elements. Bearings may be used as support for high speed vehicles such as maglev trains. Bearings are used to support rotating shafts for engines, motors, generators, fans, compressors, turbines and the like. Giant roller bearings may be used to support buildings and physical infrastructure. Different types of bearings may be used to support conventional, planetary and other types of gears. Bearings may be used to support transmissions and gear boxes such as roller thrust bearings, for example. Bearings may be used to support wheels, wheel hubs and other rolling parts using tapered roller bearings.

[0726] There are many different types of bearings such as roller bearings, needle bearings, sleeve bearings, ball bearings, radial bearings, thrust load bearings including ball thrust bearings used in

low speed applications and roller thrust bearings, taper bearings and tapered roller bearings, specialized bearings, magnetic bearings, giant roller bearings, jewel bearings (e.g., Sapphire), fluid bearings, flexure bearings to support bending element loads, and the like. References to bearings throughout this disclosure are intended to include, but not be limited by, the terms listed above.

[0727] In example embodiments, information about the health or other status or state information of or regarding a bearing in a piece of industrial equipment or in an industrial process may be obtained by monitoring the condition of various components of the industrial equipment or industrial process. Monitoring may include monitoring the amplitude and/or frequency and/or phase of a sensor signal measuring attributes such as temperature, humidity, acceleration, displacement and the like.

[0728] An embodiment of a data monitoring device 9200 is shown in Figure 75 and may include a plurality of sensors 9206 communicatively coupled to a controller 9202. The controller 9202 may include a data acquisition circuit 9204, a data storage circuit 9216, a signal evaluation circuit 9208 and, optionally, a response circuit 9210. The signal evaluation circuit 9208 may comprise a frequency transformation circuit 9212 and a frequency evaluation circuit 9214.

[0729] The plurality of sensors 9206 may be wired to ports 9226 (reference Figure 76) on the data acquisition circuit 9204. The plurality of sensors 9206 may be wirelessly connected to the data acquisition circuit 9204. The data acquisition circuit 9204 may be able to access detection values corresponding to the output of at least one of the plurality of sensors 9206 where the sensors 9206 may be capturing data on different operational aspects of a bearing or piece of equipment or infrastructure.

[0730] The selection of the plurality of sensors 9206 for a data monitoring device 9200 designed for a specific bearing or piece of equipment may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, reliability of the sensors, and the like. The impact of failure may drive the extent to which a bearing or piece of equipment is monitored with more sensors and/or higher capability sensors being dedicated to systems where unexpected or undetected bearing failure would be costly or have severe consequences.

[0731] The signal evaluation circuit 9208 may process the detection values to obtain information about a bearing being monitored. The frequency transformation circuit 9212 may transform one or more time-based detection values to frequency information. The transformation may be accomplished using techniques such as a digital Fast Fourier transform ("FFT"), Laplace transform, Z-transform, wavelet transform, other frequency domain transform, or other digital or analog signal analysis techniques, including, without limitation, complex analysis, including complex phase evolution analysis.

[0732] The frequency evaluation circuit 9214 (or frequency analysis circuit) may be structured to detect signals at frequencies of interest. Frequencies of interest may include frequencies higher than the frequency at which the equipment rotates (as measured by a tachometer, for instance), various harmonics and/or resonant frequencies associated with the equipment design and operating conditions such as multiples of shaft rotation velocities or other rotating components for the

equipment that is borne by the bearings. Changes in energy at frequencies close to the operating frequency may be an indicator of balance/imbalance in the system. Changes in energy at frequencies on the order of twice the operating frequency may be indicative of a system misalignment – for example, on the coupling, or a looseness in the system, (e.g., rattling at harmonics of the operating frequency). Changes in energy at frequencies close to three or four times the operating frequency, corresponding to the number of bolts on a coupling, may indicate wear of on one of the couplings. Changes in energy at frequencies of four, five, or more times the operating frequency may relate back to something that has a corresponding number of elements, such as if there are energy peaks or activity around five times the operating frequency there may be wear or an imbalance in a five-vane pump or the like.

[0733] In an illustrative and non-limiting example, in the analysis of roller bearings, frequencies of interest may include ball spin frequencies, cage spin frequencies, inner race frequency (as bearings often sit on a race inside a cage), outer race frequency and the like. Bearings that are damaged or beginning to fail may show humps of energy at the frequencies mentioned above and elsewhere in this disclosure. The energy at these frequencies may increase over time as the bearings wear more and become more damaged due to more variations in rotational acceleration and pings.

[0734] In an illustrative and non-limiting example, bad bearings may show humps of energy and the intensity of high frequency measurements may start to grow over time as bearings wear and become imperfect (greater acceleration and pings may show up in high frequency measurement domains). Those measurements may be indicators of air gaps in the bearing system. As bearings begin to wear, harder hits may cause the energy signal to move to higher frequencies.

[0735] In example embodiments, the signal evaluation circuit 9208 may also include one or more of a phase detection circuit, a phase lock loop circuit, a bandpass filter circuit, a peak detection circuit, and the like.

[0736] In example embodiments, the signal evaluation circuit 9208 may include a transitory signal analysis circuit. Transient signals may cause small amplitude vibrations. However, the challenge in bearing analysis is that you may receive a signal associated with a single or non-periodic impact and an exponential decay. Thus, the oscillation of the bearing may not be represented by a single sine wave, but rather by a spectrum of many high frequency sine waves. For example, a signal from a failing bearing may only be seen, in a time-based signal, as a low amplitude spike for a short amount of time. A signal from a failing bearing may be lower in amplitude than a signal associated with an imbalance even though the consequences of a failed bearing may be more significant. It is important to be able to identify these signals. This type of low amplitude, transient signal may be best analyzed using transient analysis rather than a conventional frequency transformation, such as an FFT, which would treat the signal like a low frequency sine wave. A higher resolution data stream may also provide additional data for the detection of transitory signals in low speed operations. The identification of transitory signals may enable the identification of defects in a piece of equipment or component operating at low RPMs.

[0737] In example embodiments, the transitory signal analysis circuit for bearing analysis may include envelope modulation analysis and other transitory signal analysis techniques. The signal

evaluation circuit 9208 may store along stream of detection values to the data storage circuit 9216. The transitory signal analysis circuit may use envelope analysis techniques on those long streams of detection values to identify transient effects (such as impacts) which may not be identified by conventional sine wave analysis (such as FFTs).

[0738] The signal evaluation circuit 9208 may utilize transitory signal analysis models optimized for the type of component being measured such as bearings, gears, variable speed machinery and the like. In an illustrative and non-limiting example, a gear may resonate close to its average rotational speed. In an illustrative and non-limiting example, a bearing may resonate close to the bearing rotation frequency and produce a ringing in amplitude around that frequency. For example, if the shaft inner race is wearing there may be chatter between the inner race and the shaft resulting in amplitude modulation to the left and right of the bearing frequency. The amplitude modulation may demonstrate its own sine wave characteristics with its own side bands. Various signal processing techniques may be used to eliminate the sinusoidal component, resulting in a modulation envelope for analysis.

[0739] The signal evaluation circuit 9208 may be optimized for variable speed machinery. Historically, variable speed machinery was expensive to make, and it was common to use DC motors and variable sheaves, such that flow could be controlled using vanes. Variable speed motors became more common with solid-state drive advances ("SCR devices"). The base operating frequency of equipment may be varied from the 50-60 Hz provided by standard utility companies and either slowed down or sped up to run the equipment at different speeds depending on the application. The ability to run the equipment at varying speeds may result in energy savings. However, depending on the equipment geometry, there may be some speeds which create vibrations at resonant frequencies, reducing the life of the components. Variable speed motors may also emit electricity into bearings which may damage the bearings. In example embodiments, the analysis of long data streams for envelope modulation analysis and other transitory signal analysis techniques as described herein may be useful in identifying these frequencies such that control schemes for the equipment may be designed to avoid those speeds which result in unacceptable vibrations and/or damage to the bearings.

[0740] In an illustrative and non-limiting example, heating, ventilation and air conditioning ("HVAC") systems may be assembled on site using variable speed motors, fans, belts, compressors and the like where the operating speeds are not constant, and their relative relationships are unknown. In an illustrative and non-limiting example, variable speed motors may be used in fan pumps for building air circulation. Variable speed motors may be used to vary the speed of conveyors – for example, in manufacturing assembly lines or steel mills. Variable speed motors may be used for fans in a pharmaceutical process, such as where it may be critical to avoid vibration.

[0741] In an illustrative and non-limiting example, sleeve bearings may be analyzed for defects. Sleeve bearings typically have an oil system. If the oil flow stops or the oil becomes severely contaminated, failure can occur very quickly. Therefore, a fluid particulate sensor or fluid pressure sensors may be an important source of detection values.

[0742] In an illustrative and non-limiting example, fan integrity may be evaluated by measuring air pulsations related to blade pass frequencies. For example, if a fan has 12 blades, 12 air pulsations may be measured. Variations in the amplitude of the pulsations associated with the different blades may be indicative of changes in a fan blade. Changes in frequencies associated with the air pulsations may be indicative of bearing problems.

[0743] In an illustrative and non-limiting example, compressors used in in the gas and oil field or in gas handling equipment on an assembly line may be evaluated by measuring the periodic increases in energy/pressure in the storage vessel as gas is pumped into the vessel. Periodic variations in the amplitude of the energy increases may be associated with piston wear or damage to a portion of a rotary screw. Phase evaluation of the energy signal relative to timing signals may be helpful in identifying which piston or portion of the rotary screw has damage. Changes in frequencies associated with the energy pulsations may be indicative of bearing problems.

[0744] In an illustrative and non-limiting example, cavitation/air pockets in pumps may create shuttering in the pump housing and the output flow which may be identified with the frequency transformation and frequency analysis techniques described above and elsewhere herein.

[0745] In an illustrative and non-limiting example, the frequency transformation and frequency analysis techniques described above and elsewhere herein may assist in the identification of problems in components of building HVAC systems such as big fans. If the dampers of the system are set poorly it may result in ducts pulsing or vibrating as air is pushed through the system. Monitoring of vibration sensors on the ducts may assist in the balancing of the system. If there are defects in the blades of the big fan this may also result in uneven air flow and resulting pulsation in the ductwork of the building.

[0746] In an illustrative and non-limiting example, detection values from acoustical sensors located close to the bearings may assist in the identification of issues in the engagement between gears or bad bearings. Based on a knowledge of gear ratios, such as the "in" and "out" gear ratios, for a system and measurements of the input and output rotational speed, detection values may be evaluated for energy occurring at those ratios, which in turn may be used to identify bad bearings. This could be done with simple off-the-shelf motors rather than requiring extensive retrofitting of the motor with sensors.

[0747] Based on the output of its various components, the signal evaluation circuit 9208 may make a bearing life prediction, identify a bearing health parameter, identify a bearing performance parameter, determine a bearing health parameter (e.g., fault conditions), and the like. The signal evaluation circuit 9208 may identify wear on a bearing, identify the presence of foreign matter (e.g., particulates) in the bearings, identify air gaps or a loss of fluid in oil/fluid coated bearings, identify a loss of lubrication in a set of bearings, identify a loss of power for magnetic bearings and the like, identify strain/stress of flexure bearings, and the like. The signal evaluation circuit 9208 may identify optimal operation parameters for a piece of equipment to extend bearing life. The signal evaluation circuit 9208 may identify behavior (resonant wobble) at a selected operational frequency (e.g., shaft rotation rate).

[0748] The signal evaluation circuit 9208 may communicate with the data storage circuit 9216 to access equipment specifications, equipment geometry, bearing specifications, bearing materials, anticipated state information for a plurality of bearing types, operational history, historical detection values, and the like for use in assessing the output of its various components. The signal evaluation circuit 9208 may buffer a subset of the plurality of detection values, intermediate data such as time-based detection values transformed to frequency information, filtered detection values, identified frequencies of interest, and the like for a predetermined length of time. The signal evaluation circuit 9208 may periodically store certain detection values in the data storage circuit 9216 to enable the tracking of component performance over time. In example embodiments, based on relevant operating conditions and/or failure modes that may occur as detection values approach one or more criteria, the signal evaluation circuit 9208 may store data in the data storage circuit 9216 based on the fit of data relative to one or more criteria, such as those described throughout this disclosure. Based on one sensor input meeting or approaching specified criteria or range, the signal evaluation circuit 9208 may store additional data such as RPMs, component loads, temperatures, pressures, vibrations or other sensor data of the types described throughout this disclosure in the data storage circuit 9216. The signal evaluation circuit 9208 may store data at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0749] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating and the like, sensors 9206 may comprise one or more of, without limitation, a vibration sensor, an optical vibration sensor, a thermometer, a hygrometer, a voltage sensor, a current sensor, an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, an infrared sensor, an acoustic wave sensor, a heat flux sensor, a displacement sensor, a turbidity meter, a viscosity meter, a load sensor, a tri-axial vibration sensor, an accelerometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an acoustical sensor, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference. The sensors may typically comprise at least a temperature sensor, a load sensor, a tri-axial sensor and a tachometer.

[0750] The sensors 9206 may provide a stream of data over time that has a phase component, such as relating to acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 9206 may provide a stream of data that is not conventionally phase-based, such as temperature, humidity, load, and the like. The sensors 9206 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0751] In example embodiments, as illustrated in Figure 75, the sensors 9206 may be part of the data monitoring device 9200, referred to herein in some cases as a data collector, which in some cases may comprise a mobile or portable data collector. In example embodiments, as illustrated in Figures 76 and 77, one or more external sensors 9224, which are not explicitly part of a monitoring device 9218 but rather are new, previously attached to or integrated into the equipment or component, may be opportunistically connected to or accessed by the monitoring device 9218. The monitoring device 9218 may include a controller 9220. The controller 9202 may include a data acquisition circuit 9222, a data storage circuit 9216, a signal evaluation circuit 9208 and, optionally, a response circuit 9210. The signal evaluation circuit 9208 may comprise a frequency transformation circuit 9212 and a frequency analysis circuit 9214. The data acquisition circuit 9222 may include one or more input ports 9226. The one or more external sensors 9224 may be directly connected to the one or more input ports 9226 on the data acquisition circuit 9222 of the controller 9220 or may be accessed by the data acquisition circuit 9222 wirelessly, such as by a reader, interrogator, or other wireless connection, such as over a short-distance wireless protocol. In embodiments as shown in Figure 77, a data acquisition circuit 9222 may further comprise a wireless communications circuit 9262. The data acquisition circuit 9222 may use the wireless communications circuit 9262 to access detection values corresponding to the one or more external sensors 9224 wirelessly or via a separate source or some combination of these methods.

[0752] In example embodiments, as illustrated in Figure 78, the data acquisition circuit 9222 may further comprise a multiplexer circuit 9236 as described elsewhere herein. Outputs from the multiplexer circuit 9236 may be utilized by the signal evaluation circuit 9208. The response circuit 9210 may have the ability to turn on and off portions of the multiplexor circuit 9236. The response circuit 9210 may have the ability to control the control channels of the multiplexor circuit 9236.

[0753] The response circuit 9210 may initiate actions based on a bearing performance parameter, a bearing health value, a bearing life prediction parameter, and the like. The response circuit 9210 may evaluate the results of the signal evaluation circuit 9208 and, based on certain criteria or the output from various components of the signal evaluation circuit 9208, initiate an action. The criteria may include a sensor's detection values at certain frequencies or phases relative to a timer signal where the frequencies or phases of interest may be based on the equipment geometry, equipment control schemes, system input, historical data, current operating conditions, and/or an anticipated response. The criteria may include a sensor's detection values at certain frequencies or phases relative to detection values of a second sensor. The criteria may include signal strength at certain resonant frequencies/harmonics relative to detection values associated with a system tachometer or anticipated based on equipment geometry and operation conditions. Criteria may include a predetermined peak value for a detection value from a specific sensor, a cumulative value of a sensor's corresponding detection value over time, a change in peak value, a rate of change in a peak value, and/or an accumulated value (e.g., a time spent above/below a threshold value, a weighted time spent above/below one or more threshold values, and/or an area of the detected value above/below one or more threshold values). The criteria may comprise combinations of data from different sensors such as relative values, relative changes in value, relative rates of change in

value, relative values over time, and the like. The relative criteria may change with other data or information such as process stage, type of product being processed, type of equipment, ambient temperature and humidity, external vibrations from other equipment, and the like. The relative criteria may be reflected in one or more calculated statistics or metrics (including ones generated by further calculations on multiple criteria or statistics), which in turn may be used for processing (such as on-board a data collector or by an external system), such as to be provided as an input to one or more of the machine learning capabilities described in this disclosure, to a control system (which may be on board a data collector or remote, such as to control selection of data inputs, multiplexing of sensor data, storage, or the like), or as a data element that is an input to another system, such as a data stream or data package that may be available to a data marketplace, a SCADA system, a remote control system, a maintenance system, an analytic system, or other system.

[0754] Certain embodiments are described herein as detected values exceeding thresholds or predetermined values, but detected values may also fall below thresholds or predetermined values – for example, where an amount of change in the detected value is expected to occur, but detected values indicate that the change may not have occurred. For example, and without limitation, vibrational data may indicate system agitation levels, properly operating equipment, or the like, and vibrational data below amplitude and/or frequency thresholds may be an indication of a process that is not operating according to expectations. Except where the context clearly indicates otherwise, any description herein describing a determination of a value above a threshold and/or exceeding a predetermined or expected value is understood to include determination of a value below a threshold and/or falling below a predetermined or expected value.

[0755] The predetermined acceptable range may be based on anticipated system response or vibration based on the equipment geometry and control scheme such as number of bearings, relative rotational speed, influx of power to the system at a certain frequency, and the like. The predetermined acceptable range may also be based on long term analysis of detection values across a plurality of similar equipment and components and correlation of data with equipment failure.

[0756] In example embodiments, an alert may be issued based on some of the criteria discussed above. In an illustrative example, an increase in temperature and energy at certain frequencies may indicate a hot bearing that is starting to fail. In example embodiments, the relative criteria for an alarm may change with other data or information such as process stage, type of product being processed on equipment, ambient temperature and humidity, external vibrations from other equipment and the like. In an illustrative and non-limiting example, the response circuit 9210 may initiate an alert if a vibrational amplitude and/or frequency exceeds a predetermined maximum value, if there is a change or rate of change that exceeds a predetermined acceptable range, and/or if an accumulated value based on vibrational amplitude and/or frequency exceeds a threshold.

[0757] In example embodiments, response circuit 9210 may cause the data acquisition circuit 9204 to enable or disable the processing of detection values corresponding to certain sensors based on some of the criteria discussed above. This may include switching to sensors having different response rates, sensitivity, ranges, and the like, or accessing new sensors or types of sensors, and

the like. Switching may be undertaken based on a model, a set of rules, or the like. In example embodiments, switching may be under control of a machine learning system, such that switching is controlled based on one or more metrics of success, combined with input data, over a set of trials, which may occur under supervision of a human supervisor or under control of an automated system. Switching may involve switching from one input port to another (such as to switch from one sensor to another). Switching may involve altering the multiplexing of data, such as combining different streams under different circumstances. Switching may also involve activating a system to obtain additional data, such as moving a mobile system (such as a robotic or drone system), to a location where different or additional data is available (such as positioning an image sensor for a different view or positioning a sonar sensor for a different direction of collection) or to a location where different sensors can be accessed (such as moving a collector to connect up to a sensor that is disposed at a location in an environment by a wired or wireless connection). This switching may be implemented by changing the control signals for a multiplexor circuit 9236 and/or by turning on or off certain input sections of the multiplexor circuit 9236. The response circuit 9210 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 9210 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0758] In example embodiments, the response circuit 9210 may recommend maintenance at an upcoming process stop or initiate a maintenance call. The response circuit 9210 may recommend changes in process or operating parameters to remotely balance the piece of equipment. In example embodiments, the response circuit 9210 may implement or recommend process changes – for example to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but still at least partially operational, to change the operating speed of a component (such as to put it in a lower-demand mode), to initiate amelioration of an issue (such as to signal for additional lubrication of a roller bearing set, or to signal for an alignment process for a system that is out of balance), and the like.

[0759] In embodiments as shown in Figures 79, 80, 81, and 82, a data monitoring system 9240 may include at least one data monitoring device 9250. The at least one data monitoring device 9250 may include sensors 9206 and a controller 9242 comprising a data acquisition circuit 9204, a signal evaluation circuit 9208, a data storage circuit 9216, and a communications circuit 9246. The signal evaluation circuit 9208 may include at least one of a frequency detection circuit 9212 and a frequency analysis circuit 9214. There may also be an optional response circuit as described above and elsewhere herein. The signal evaluation circuit 9208 may periodically share data with the communication circuit 9246 for transmittal to a remote server 9244 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 9248. Because relevant operating conditions and/or failure modes may occur as sensor values approach one or more criteria, the signal evaluation circuit 9208 may share data with the communication circuit 9246 for transmittal to the remote server 9244 based on the fit of data relative to one or more criteria. Based on one sensor input meeting or approaching specified criteria

or range, the signal evaluation circuit 9208 may share additional data such as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The signal evaluation circuit 9208 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server.

[0760] In example embodiments, as shown in Figure 79, the communications circuit 9246 may communicate data directly to a remote server 9244. In example embodiments, as shown in Figure 80, the communications circuit 9246 may communicate data to an intermediate computer 9252, which may include a processor 9254 running an operating system 9256 and a data storage circuit 9258. The intermediate computer 9252 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 9244.

[0761] In example embodiments, as illustrated in Figures 81 and 82, a data collection system 9260 may have a plurality of monitoring devices 9250 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment, (both the same and different types of equipment) in the same facility as well as collecting data from monitoring devices in multiple facilities. A monitoring application 9248 on a remote server 9244 may receive and store one or more of the following: detection values, timing signals and data coming from a plurality of the various monitoring devices 9250. In example embodiments, as shown in Figure 81, the communications circuit 9246 may communicate data directly to a remote server 9244. In example embodiments, as shown in Figure 82, the communications circuit 9246 may communicate data to an intermediate computer 9252, which may include a processor 9254 running an operating system 9256 and a data storage circuit 9258. There may be an individual intermediate computer 9252 associated with each monitoring device 9264 or an individual intermediate computer 9252 may be associated with a plurality of monitoring devices 9250 where the intermediate computer 9252 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 9244.

[0762] The monitoring application 9248 may select subsets of the detection values, timing signals and data to be jointly analyzed. Subsets for analysis may be selected based on a bearing type, bearing materials, or a single type of equipment in which a bearing is operating. Subsets for analysis may be selected or grouped based on common operating conditions or operational history such as size of load, operational condition (e.g., intermittent, continuous), operating speed or tachometer, common ambient environmental conditions such as humidity, temperature, air or fluid particulate, and the like. Subsets for analysis may be selected based on common anticipated state information. Subsets for analysis may be selected based on the effects of other nearby equipment such as nearby machines rotating at similar frequencies, nearby equipment producing electromagnetic fields, nearby equipment producing heat, nearby equipment inducing movement or vibration, nearby equipment emitting vapors, chemicals or particulates, or other potentially interfering or intervening effects.

[0763] The monitoring application 9248 may analyze a selected subset. In an illustrative example, data from a single component may be analyzed over different time periods, such as one operating cycle, cycle-to-cycle comparisons, trends over several operating cycles/times such as a month, a

year, the life of the component, or the like. Data from multiple components of the same type may also be analyzed over different time periods. Trends in the data such as changes in frequency or amplitude may be correlated with failure and maintenance records associated with the same component or piece of equipment. Trends in the data such as changing rates of change associated with start-up or different points in the process may be identified. Additional data may be introduced into the analysis such as output product quality, output quantity (such as per unit of time), indicated success or failure of a process, and the like. Correlation of trends and values for different types of data may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected performance. The analysis may identify model improvements to the model for anticipated state information, recommendations around sensors to be used, positioning of sensors and the like. The analysis may identify additional data to collect and store. The analysis may identify recommendations regarding needed maintenance and repair and/or the scheduling of preventative maintenance. The analysis may identify recommendations around purchasing replacement bearings and the timing of the replacement of the bearings. The analysis may result in warning regarding the dangers of catastrophic failure conditions. This information may be transmitted back to the monitoring device to update types of data collected and analyzed locally or to influence the design of future monitoring devices.

[0764] In example embodiments, the monitoring application 9248 may have access to equipment specifications, equipment geometry, bearing specifications, bearing materials, anticipated state information for a plurality of bearing types, operational history, historical detection values, bearing life models and the like for use analyzing the selected subset using rule-based or model-based analysis. In example embodiments, the monitoring application 9248 may feed a neural net with the selected subset to learn to recognize various operating state, health states (e.g., lifetime predictions) and fault states utilizing deep learning techniques. In example embodiments, a hybrid of the two techniques (model-based learning and deep learning) may be used.

[0765] In an illustrative and non-limiting example, the health of bearings on conveyors and lifters in an assembly line, in water pumps on industrial vehicles and in compressors in gas handling systems, in compressors situated out in the gas and oil fields, in factory air conditioning units and in factory mineral pumps may be monitored using the frequency transformation and frequency analysis techniques, data monitoring devices and data collection systems described herein.

[0766] In an illustrative and non-limiting example, the health of one or more of bearings, gears, blades, screws and associated shafts, motors, rotors, stators, gears, and other components of gear boxes, motors, pumps, vibrating conveyors, mixers, centrifuges, drilling machines, screw drivers and refining tanks situated in the oil and gas fields may be evaluated using the frequency transformation and frequency analysis techniques, data monitoring devices and data collection systems described herein.

[0767] In an illustrative and non-limiting example, the health of bearings and associated shafts, motors, rotors, stators, gears, and other components of rotating tank/mixer agitators, mechanical/rotating agitators, and propeller agitators, to promote chemical reactions deployed in chemical and pharmaceutical production lines may be evaluated using the frequency

transformation and frequency analysis techniques, data monitoring devices and data collection systems described herein.

[0768] In an illustrative and non-limiting example, the health of bearings and associated shafts, motors, rotors, stators, gears, and other components of vehicle systems such as steering mechanisms or engines may be evaluated using the frequency transformation and frequency analysis techniques, data monitoring devices and data collection systems described herein.

[0769] An example monitoring device for bearing analysis in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; and a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter.

[0770] In certain further embodiments, an example monitoring device includes one or more of: a response circuit to perform at least one operation in response to the bearing performance parameter, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer; wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference.

[0771] An example monitoring device for bearing analysis in an industrial environment, the monitoring device includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; and a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing health value.

[0772] In certain example embodiments, an example monitoring device further includes one or more of: a response circuit to perform at least one operation in response to the bearing health value, wherein the plurality of input sensors includes at least two sensors selected from the group

consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer; wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference.

[0773] An example monitoring device for bearing analysis in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; and a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing life prediction parameter.

[0774] In certain example embodiments, a monitoring device further includes one or more of: a response circuit to perform at least one operation in response to the bearing life prediction parameter, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer; wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference.

[0775] An example monitoring device for bearing analysis in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a

predetermined length of time; and a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter, wherein the data acquisition circuit comprises a multiplexer circuit whereby alternative combinations of the detection values may be selected based on at least one of user input, a detected state and a selected operating parameter for a machine.

[0776] In certain further embodiments, an example monitoring device further includes one or more of: a response circuit to perform at least one operation in response to the bearing performance parameter, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer; a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference; wherein the at least one operation comprises enabling or disabling one or more portions of the multiplexer circuit, or altering the multiplexer control lines; wherein the data acquisition circuit comprises at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits.

[0777] An example system for data collection, processing, and bearing analysis in an industrial environment includes: a plurality of monitoring devices, each monitoring device comprising: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing life prediction; a communication circuit structured to communicate with a remote server providing the bearing life prediction and a portion of the buffered detection values to the remote server; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the plurality of monitoring devices.

[0778] In certain further embodiments, an example monitoring device includes one or more of: a response circuit to perform at least one operation in response to the bearing life prediction, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer; wherein the at least one operation is further in response to at least one of: a change in amplitude of at least

one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one of the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one of the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein the storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference.

[0779] An example system for data collection, processing, and bearing analysis in an industrial environment comprising: a plurality of monitoring devices, each comprising: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing specifications and anticipated state information for a plurality of bearing types and buffering the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; a communication circuit structured to communicate with a remote server providing the life prediction and a portion of the buffered detection values to the remote server; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the plurality of monitoring devices.

[0780] In certain further embodiments, an example monitoring device further includes one or more of: a response circuit to perform at least one operation in response to the bearing performance parameter, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer; wherein the at least one operation is further in response to at least one of: a change in amplitude of at least one of the plurality of detection values; a change in frequency or relative phase of at least one of the plurality of detection values; a rate of change in both amplitude and relative phase of at least one the plurality of detection values; and a relative rate of change in amplitude and relative phase of at least one the plurality of detection values; wherein the at least one operation comprises issuing an alert; wherein the alert may be one of haptic, audible and visual; wherein the at least one operation further comprises storing additional data in the data storage circuit; wherein storing additional data in the data storage circuit is further in response to at least one of: a change in the relative phase difference and a relative rate of change in the relative phase difference.

[0781] An example system for data collection, processing, and bearing analysis in an industrial environment includes: a plurality of monitoring devices, each monitoring device comprising: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a streaming circuit for streaming at least a subset of the

acquired detection values to a remote learning system; and a remote learning system including a bearing analysis circuit structured to analyze the detection values relative to a machine-based understanding of the state of the at least one bearing.

[0782] In certain further embodiments, an example system further includes one or more of: wherein the machine-based understanding is developed based on a model of the bearing that determines a state of the at least one bearing based at least in part on the relationship of the behavior of the bearing to an operating frequency of a component of the industrial machine; wherein the state of the at least one bearing is at least one of an operating state, a health state, a predicted lifetime state and a fault state; wherein the machine-based understanding is developed based by providing inputs to a deep learning machine, wherein the inputs comprise a plurality of streams of detection values for a plurality of bearings and a plurality of measured state values for the plurality of bearings; wherein the state of the at least one bearing is at least one of an operating state, a health state, a predicted lifetime state and a fault state.

[0783] An example method of analyzing bearings and sets of bearings, includes: receiving a plurality of detection values corresponding to data from a temperature sensor, a vibration sensor positioned near the bearing or set of bearings and a tachometer to measure rotation of a shaft associated with the bearing or set of bearings; comparing the detection values corresponding to the temperature sensor to a predetermined maximum level; filtering the detection values corresponding to the vibration sensor through a high pass filter where the filter is selected to eliminate vibrations associated with detection values associated with the tachometer; identifying rapid changes in at least one of a temperature peak and a vibration peak; identifying frequencies at which spikes in the filtered detection values corresponding to the vibration sensor occur and comparing frequencies and spikes in amplitude relative to an anticipated state information and specification associated with the bearing or set of bearings; and determining a bearing health parameter.

[0784] An example device for monitoring roller bearings in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage circuit structured to store specifications and anticipated state information for a plurality of types of roller bearings and buffering the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a response circuit to perform at least one operation in response to the bearing performance prediction, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer.

[0785] An example device for monitoring sleeve bearings in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing sleeve bearing specifications and

anticipated state information for types of sleeve bearings and buffering the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a response circuit to perform at least one operation in response to the bearing performance parameter, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer.

[0786] An example system for monitoring pump bearings in an industrial environment, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing pump specifications, bearing specifications, anticipated state information for pump bearings and buffering the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a response circuit to perform at least one operation in response to the bearing performance parameter, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor and a tachometer.

[0787] An example system for collection, processing, and analyzing pump bearings in an industrial environment includes: a plurality of monitoring devices, each comprising: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors communicatively coupled to the data acquisition circuit; a data storage for storing pump specifications, bearing specifications, anticipated state information for pump bearings and buffering the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to the pump and bearing specifications and anticipated state information resulting in a bearing performance parameter; a communication circuit structured to communicate with a remote server providing the bearing performance parameter and a portion of the buffered detection values to the remote server; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the plurality of monitoring devices.

[0788] An example system for estimating a conveyor health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the conveyor and associated rotating components, store historical conveyor and component performance and buffer the plurality of detection values for a predetermined length of time; a bearing analysis circuit

structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a system analysis circuit structured to utilize the bearing performance and at least one of an anticipated state, historical data and a system geometry to estimate a conveyor health performance.

[0789] An example system for estimating an agitator health parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the agitator and associated components, store historical agitator and component performance and buffer the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a system analysis circuit structured to utilize the bearing performance and at least one of an anticipated state, historical data and a system geometry to estimate an agitation health parameter. In certain further embodiments, an example device further includes where the agitator is one of a rotating tank mixer, a large tank mixer, a portable tank mixers, a tote tank mixer, a drum mixer, a mounted mixer and a propeller mixer.

[0790] An example system for estimating a vehicle steering system performance parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the vehicle steering system, the rack, the pinion, and the steering column, store historical steering system performance and buffer the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a system analysis circuit structured to utilize the bearing performance and at least one of an anticipated state, historical data and a system geometry to estimate a vehicle steering system performance parameter.

[0791] An example system for estimating a pump performance parameter, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the pump and pump components, store historical steering system performance and buffer the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information

resulting in a bearing performance parameter; a system analysis circuit structured to utilize the bearing performance and at least one of an anticipated state, historical data and a system geometry to estimate a pump performance parameter. In certain example embodiments, and example system further includes wherein the pump is a water pump in a car, and/or wherein the pump is a mineral pump.

[0792] An example system for estimating a performance parameter for a drilling machine, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the drilling machine and drilling machine components, store historical drilling machine performance and buffer the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a system analysis circuit structured to utilize the bearing performance and at least one of an anticipated state, historical data and a system geometry to estimate a performance parameter for the drilling machine. In certain further embodiments, the drilling machine is one of an oil drilling machine and a gas drilling machine.

[0793] An example system for estimating a performance parameter for a drilling machine, includes: a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for the drilling machine and drilling machine components, store historical drilling machine performance and buffer the plurality of detection values for a predetermined length of time; a bearing analysis circuit structured to analyze buffered detection values relative to specifications and anticipated state information resulting in a bearing performance parameter; and a system analysis circuit structured to utilize bearing performance and at least one of an anticipated state, historical data and a system geometry to estimate a performance parameter for the drilling machine.

[0794] Rotating components are used throughout many different types of equipment and applications. Rotating components may include shafts, motors, rotors, stators, bearings, fins, vanes, wings, blades, fans, bearings, wheels, hubs, spokes, balls, rollers, pins, gears and the like. In example embodiments, information about the health or other status or state information of or regarding a rotating component in a piece of industrial equipment or in an industrial process may be obtained by monitoring the condition of the component or various other components of the industrial equipment or industrial process and identifying torsion on the component. Monitoring

may include monitoring the amplitude and phase of a sensor signal, such as one measuring attributes such as angular position, angular velocity, angular acceleration, and the like.

[0795] An embodiment of a data monitoring device 9400 is shown in Figure 83 and may include a plurality of sensors 9406 communicatively coupled to a controller 9402. The controller 9402 may include a data acquisition circuit 9404, a data storage circuit 9414, a system evaluation circuit 9408 and, optionally, a response circuit 9410. The system evaluation circuit 9408 may comprise a torsion analysis circuit 9412.

[0796] The plurality of sensors 9406 may be wired to ports on the data acquisition circuit 9404. The plurality of sensors 9406 may be wirelessly connected to the data acquisition circuit 9404. The data acquisition circuit 9404 may be able to access detection values corresponding to the output of at least one of the plurality of sensors 9406 where the sensors 9406 may be capturing data on different operational aspects of a bearing or piece of equipment or infrastructure.

[0797] The selection of the plurality of sensors 9406 for a data monitoring device 9400 designed to assess torsion on a component, such as a shaft, motor, rotor, stator, bearing or gear, or other component described herein, or a combination of components, such as within or comprising a drive train or piece of equipment or system, may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, reliability of the sensors, and the like. The impact of failure may drive the extent to which a bearing or piece of equipment is monitored with more sensors and/or higher capability sensors being dedicated to systems where unexpected or undetected bearing failure would be costly or have severe consequences. To assess torsion the sensors may include, among other options, an angular position sensor and/or an angular velocity sensor and/or an angular acceleration sensor.

[0798] The system evaluation circuit 9408 may process the detection values to obtain information about one or more rotating components being monitored. The torsional analysis circuit 9412 may be structured to identify torsion in a component or system, such as based on anticipated state, historical state, system geometry and the like, such as that which is available from the data storage circuit 9414. The torsional analysis circuit 9412 may be structured to identify torsion using a variety of techniques such as amplitude, phase and frequency differences in the detection values from two linear accelerometers positioned at different locations on a shaft. The torsional analysis circuit 9412 may identify torsion using the difference in amplitude and phase between an angular accelerometer on a shaft and an angular accelerometer on a slip ring on the end of the shaft. The torsional analysis circuit 9412 may identify shear stress/elongation on a component using two strain gauges in a half bridge configuration or four strain gauges in a full bridge configuration. The torsional analysis circuit 9412 may use coder based techniques such as markers to identify the rotation of a shaft, bearing, rotor, stator, gear or other rotating component. The markers being assessed may include visual markers such as gear teeth or stripes on a shaft captured by an image sensor, light detector or the like. The markers being assessed may include magnetic components located on the rotating component and sensed by an electromagnetic pickup. The sensor may be a Hall Effect sensor.

[0799] Additional input sensors may include a thermometer, a heat flux sensor, a magnetometer, an axial load sensor, a radial load sensor, an accelerometer, a shear-stress torque sensor, a twist angle sensor and the like. Twist angle may include rotational information at two positions on shaft or an angular velocity or angular acceleration at two positions on a shaft. In example embodiments, the sensors may be positioned at different ends of the shaft.

[0800] The torsional analysis circuit 9412 may include one or more of a transient signal analysis circuit and/or a frequency transformation circuit and/or a frequency analysis circuit as described elsewhere herein.

[0801] In example embodiments, the transitory signal analysis circuit for torsional analysis may include envelope modulation analysis, and other transitory signal analysis techniques. The system evaluation circuit 9408 may store long stream of detection values to the data storage circuit 9414. The transitory signal analysis circuit may use envelope analysis techniques on those long streams of detection values to identify transient effects (such as impacts) which may not be identified by conventional sine wave analysis (such as FFTs).

[0802] In example embodiments, the frequencies of interest may include identifying energy at relation-order bandwidths for rotating equipment. The maximum order observed may comprise a function of the bandwidth of the system and the rotational speed of the component. For varying speeds (run-ups, run-downs, etc.), the minimum RPM may determine the maximum-observed order. In example embodiments, there may be torsional resonance at harmonics of the forcing frequency/frequency at which a component is being driven.

[0803] In an illustrative and non-limiting example, the monitoring device may be used to collect and process sensor data to measure torsion on a component. The monitoring device may be in communication with or include a high resolution, high speed vibration sensor to collect data over an extended period of time, enough to measure multiple cycles of rotation. For gear driven equipment, the sampling resolution should be such that the number of samples taken per cycle is at least equal to the number of gear teeth driving the component. It will be understood that a lower sampling resolution may also be utilized, which may result in a lower confidence determination and/or taking data over a longer period of time to develop sufficient statistical confidence. This data may then be used in the generation of a phase reference (relative probe) or tachometer signal for a piece of equipment. This phase reference may be used to align phase data such as velocity and/or positional and/or acceleration data from multiple sensors located at different positions on a component or on different components within a system. This information may facilitate the determination of torsion for different components or the generation of an Operational Deflection Shape ("ODS"), indicating the extent of torsion on one or more components during an operational mode.

[0804] The higher resolution data stream may provide additional data for the detection of transitory signals in low speed operations. The identification of transitory signals may enable the identification of defects in a piece of equipment or component.

[0805] In an illustrative and non-limiting example, the monitoring device may be used to identify mechanical jitter for use in failure prediction models. The monitoring device may begin acquiring

data when the piece of equipment starts up through ramping up to operating speed or during operation. Once at operating speed, it is anticipated that the torsional jitter should be minimal and changes in torsion during this phase may be indicative of cracks, bearing faults and the like. Additionally, known torsions may be removed from the signal to facilitate the identification of unanticipated torsions resulting from system design flaws or component wear. Having phase information associated with the data collected at operating speed may facilitate identification of a location of vibration and potential component wear. Relative phase information for a plurality of sensors located throughout a machine may facilitate the evaluation of torsion as it is propagated through a piece of equipment.

[0806] Based on the output of its various components, the system evaluation circuit 9408 may make a component life prediction, identify a component health parameter, identify a component performance parameter, and the like. The system evaluation circuit 9408 may identify unexpected torsion on a rotating component, identify strain/stress of flexure bearings, and the like. The system evaluation circuit 9408 may identify optimal operation parameters for a piece of equipment to reduce torsion and extend component life. The system evaluation circuit 9408 may identify torsion at selected operational frequencies (e.g., shaft rotation rates). Information about operational frequencies causing torsion may facilitate equipment operational balance in the future.

[0807] The system evaluation circuit 9408 may communicate with the data storage circuit 9414 to access equipment specifications, equipment geometry, bearing specifications, component materials, anticipated state information for a plurality of component types, operational history, historical detection values, and the like for use in assessing the output of its various components. The system evaluation circuit 9408 may buffer a subset of the plurality of detection values, intermediate data such as time-based detection values, time-based detection values transformed to frequency information, filtered detection values, identified frequencies of interest, and the like for a predetermined length of time. The system evaluation circuit 9408 may periodically store certain detection values in the data storage circuit 9414 to enable the tracking of component performance over time. In example embodiments, based on relevant operating conditions and/or failure modes, which may occur as detection values approach one or more criteria, the system evaluation circuit 9408 may store data in the data storage circuit 9414 based on the fit of data relative to one or more criteria, such as those described throughout this disclosure. Based on one sensor input meeting or approaching specified criteria or range, the system evaluation circuit 9408 may store additional data such as RPM information, component loads, temperatures, pressures, vibrations or other sensor data of the types described throughout this disclosure in the data storage circuit 9414. The system evaluation circuit 9408 may store data in the data storage circuit at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0808] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating and the like, sensors 9406 may comprise, without limitation, one or more of the following: a displacement sensor, an angular velocity sensor, an angular

accelerometer, a vibration sensor, an optical vibration sensor, a thermometer, a hygrometer, a voltage sensor, a current sensor, an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, an infrared sensor, an acoustic wave sensor, a heat flux sensor, a displacement sensor, a turbidity meter, a viscosity meter, a load sensor, a tri-axial vibration sensor, an accelerometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an acoustical sensor, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference.

[0809] The sensors 9406 may provide a stream of data over time that has a phase component, such as relating to angular velocity, angular acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 9406 may provide a stream of data that is not conventionally phase-based, such as temperature, humidity, load, and the like. The sensors 9406 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0810] In an illustrative and non-limiting example, when assessing engine components it may be desirable to remove vibrations due to the timing of piston vibrations or anticipated vibrational input due to crankshaft geometry to assist in identifying other torsional forces on a component. This may assist in assessing the health of such diverse components as a water pump in a vehicle or positive displacement pumps.

[0811] In an illustrative and non-limiting example, torsional analysis and the identification of variations in torsion may assist in the identification of stick-slip in a gear or transfer system. In some cases, this may only occur once per cycle, and phase information may be as important as or more important than the amplitude of the signal in determining system state or behavior.

[0812] In an illustrative and non-limiting example, torsional analysis may assist in the identification, prediction (e.g., timing) and evaluation of lash in a drive train and the follow-on torsion resulting from a change in direction or start up, which in turn may be used for controlling a system, assessing needs for maintenance, assessing needs for balancing or otherwise re-setting components, or the like.

[0813] In an illustrative and non-limiting example, when assessing compressors, it may be desirable to remove vibrations due to the timing of piston vibrations or anticipated vibrational input associated with the techniques and geometry used for positive displacement compressors to assist in identifying other torsional forces on a component. This may assist in assessing the health of compressors in such diverse environments as air conditioning units in factories, compressors in gas handling systems in an industrial environment, compressors in oil fields, and other environments as described elsewhere herein.

[0814] In an illustrative and non-limiting example, torsional analysis may facilitate the understanding of the health and expected life of various components associated with the drive trains of vehicles, such as cranes, bulldozers, tractors, haulers, backhoes, forklifts, agricultural

equipment, mining equipment, boring and drilling machines, digging machines, lifting machines, mixers (e.g., cement mixers), tank trucks, refrigeration trucks, security vehicles (e.g., including safes and similar facilities for preserving valuables), underwater vehicles, watercraft, aircraft, automobiles, trucks, trains and the like, as well as drive trains of moving apparatus, such as assembly lines, lifts, cranes, conveyors, hauling systems, and others. The evaluation of the sensor data with the model of the system geometry and operating conditions may be useful in identifying unexpected torsion and the transmission of that torsion from the motor and drive shaft, from the drive shaft to the universal joint and from the universal joint to one or more wheel axles.

[0815] In an illustrative and non-limiting example, torsional analysis may facilitate in the understanding of the health and expected life of various components associated with train/tram wheels and wheel sets. As discussed above, torsional analysis may facilitate in the identification of stick-slip between the wheels or wheel sets and the rail. The torsional analysis in view of the system geometry may facilitate the identification of torsional vibration due to stick-slip as opposed to the torsional vibration due to the driving geometry connecting the engine to the drive shaft to the wheel axle.

[0816] In example embodiments, as illustrated in Figure 83, the sensors 9406 may be part of the data monitoring device 9400, referred to herein in some cases as a data collector, which in some cases may comprise a mobile or portable data collector. In example embodiments, as illustrated in Figures 84 and 85, one or more external sensors 9422, which are not explicitly part of a monitoring device 9416 but rather are new, previously attached to or integrated into the equipment or component, may be opportunistically connected to or accessed by the monitoring device 9416. The monitoring device 9416 may include a controller 9418. The controller 9418 may include a data acquisition circuit 9420, a data storage circuit 9414, a system evaluation circuit 9408 and, optionally, a response circuit 9410. The system evaluation circuit 9408 may comprise a torsional analysis circuit 9412. The data acquisition circuit 9420 may include one or more input ports 9424. In embodiments as shown in Figure 85, a data acquisition circuit 9420 may further comprise a wireless communications circuit 9426. The one or more external sensors 9422 may be directly connected to the one or more input ports 9424 on the data acquisition circuit 9420 of the controller 9418 or may be accessed by the data acquisition circuit 9420 wirelessly using the wireless communications circuit 9426, such as by a reader, interrogator, or other wireless connection, such as over a short-distance wireless protocol. The data acquisition circuit 9420 may use the wireless communications circuit 9426 to access detection values corresponding to the one or more external sensors 9422 wirelessly or via a separate source or some combination of these methods.

[0817] In example embodiments, as illustrated in Figure 86, the data acquisition circuit 9432 may further comprise a multiplexer circuit 9434 as described elsewhere herein. Outputs from the multiplexer circuit 9434 may be utilized by the system evaluation circuit 9408. The response circuit 9410 may have the ability to turn on or off portions of the multiplexor circuit 9434. The response circuit 9410 may have the ability to control the control channels of the multiplexor circuit 9434.

[0818] The response circuit 9410 may initiate actions based on a component performance parameter, a component health value, a component life prediction parameter, and the like. The

response circuit 9410 may evaluate the results of the system evaluation circuit 9408 and, based on certain criteria or the output from various components of the system evaluation circuit 9408, may initiate an action. The criteria may include identification of torsion on a component by the torsional analysis circuit. The criteria may include a sensor's detection values at certain frequencies or phases relative to a timer signal where the frequencies or phases of interest may be based on the equipment geometry, equipment control schemes, system input, historical data, current operating conditions, and/or an anticipated response. The criteria may include a sensor's detection values at certain frequencies or phases relative to detection values of a second sensor. The criteria may include signal strength at certain resonant frequencies/harmonics relative to detection values associated with a system tachometer or anticipated based on equipment geometry and operation conditions. Criteria may include a predetermined peak value for a detection value from a specific sensor, a cumulative value of a sensor's corresponding detection value over time, a change in peak value, a rate of change in a peak value, and/or an accumulated value (e.g., a time spent above/below a threshold value, a weighted time spent above/below one or more threshold values, and/or an area of the detected value above/below one or more threshold values). The criteria may comprise combinations of data from different sensors such as relative values, relative changes in value, relative rates of change in value, relative values over time, and the like. The relative criteria may change with other data or information such as process stage, type of product being processed, type of equipment, ambient temperature and humidity, external vibrations from other equipment, and the like. The relative criteria may be reflected in one or more calculated statistics or metrics (including ones generated by further calculations on multiple criteria or statistics), which in turn may be used for processing (such as on board a data collector or by an external system), such as to be provided as an input to one or more of the machine learning capabilities described in this disclosure, to a control system (which may be on board a data collector or remote, such as to control selection of data inputs, multiplexing of sensor data, storage, or the like), or as a data element that is an input to another system, such as a data stream or data package that may be available to a data marketplace, a SCADA system, a remote control system, a maintenance system, an analytic system, or other system.

[0819] Certain embodiments are described herein as detected values exceeding thresholds or predetermined values, but detected values may also fall below thresholds or predetermined values – for example where an amount of change in the detected value is expected to occur, but detected values indicate that the change may not have occurred. Except where the context clearly indicates otherwise, any description herein describing a determination of a value above a threshold and/or exceeding a predetermined or expected value is understood to include determination of a value below a threshold and/or falling below a predetermined or expected value.

[0820] The predetermined acceptable range may be based on anticipated torsion based on equipment geometry, the geometry of a transfer system, an equipment configuration or control scheme, such as a piston firing sequence, and the like. The predetermined acceptable range may also be based on historical performance or predicted performance, such as long term analysis of signals and performance both from the past run and from the past several runs. The predetermined

acceptable range may also be based on historical performance or predicted performance, or based on long term analysis of signals and performance across a plurality of similar equipment and components (both within a specific environment, within an individual company, within multiple companies in the same industry and across industries). The predetermined acceptable range may also be based on a correlation of sensor data with actual equipment and component performance.

[0821] In example embodiments, an alert may be issued based on some of the criteria discussed above. In example embodiments, the relative criteria for an alarm may change with other data or information, such as process stage, type of product being processed on equipment, ambient temperature and humidity, external vibrations from other equipment and the like. In an illustrative and non-limiting example, the response circuit 9410 may initiate an alert if a torsion in a component across a plurality of components exceeds a predetermined maximum value, if there is a change or rate of change that exceeds a predetermined acceptable range, and/or if an accumulated value based on torsion amplitude and/or frequency exceeds a threshold.

[0822] In example embodiments, response circuit 9410 may cause the data acquisition circuit 9432 to enable or disable the processing of detection values corresponding to certain sensors based on some of the criteria discussed above. This may include switching to sensors having different response rates, sensitivity, ranges, and the like; accessing new sensors or types of sensors, and the like. Switching may be undertaken based on a model, a set of rules, or the like. In example embodiments, switching may be under control of a machine learning system, such that switching is controlled based on one or more metrics of success, combined with input data, over a set of trials, which may occur under supervision of a human supervisor or under control of an automated system. Switching may involve switching from one input port to another (such as to switch from one sensor to another). Switching may involve altering the multiplexing of data, such as combining different streams under different circumstances. Switching may involve activating a system to obtain additional data, such as moving a mobile system (such as a robotic or drone system), to a location where different or additional data is available (such as positioning an image sensor for a different view or positioning a sonar sensor for a different direction of collection) or to a location where different sensors can be accessed (such as moving a collector to connect up to a sensor that is disposed at a location in an environment by a wired or wireless connection). This switching may be implemented by changing the control signals for a multiplexor circuit 9434 and/or by turning on or off certain input sections of the multiplexor circuit 9434.

[0823] The response circuit 9410 may calculate transmission effectiveness based on differences between a measured and theoretical angular position and velocity of an output shaft after accounting for the gear ration and any phase differential between input and output.

[0824] The response circuit 9410 may identify equipment or components that are due for maintenance. The response circuit 9410 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 9410 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0825] In example embodiments, the response circuit 9410 may recommend maintenance at an upcoming process stop or initiate a maintenance call. The response circuit 9410 may recommend changes in process or operating parameters to remotely balance the piece of equipment. In example embodiments, the response circuit 9410 may implement or recommend process changes – for example to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but still at least partially operational, to change the operating speed of a component (such as to put it in a lower-demand mode), to initiate amelioration of an issue (such as to signal for additional lubrication of a roller bearing set, or to signal for an alignment process for a system that is out of balance), and the like.

[0826] In embodiments as shown in Figures 87, 88, 89, and 90, a data monitoring system 9460 may include at least one data monitoring device 9448. At least one data monitoring device 9448 may include sensors 9406 and a controller 9438 comprising a data acquisition circuit 9404, a system evaluation circuit 9408, a data storage circuit 9414, and a communications circuit 9442. The system evaluation circuit 9408 may include a torsional analysis circuit 9412. There may also be an optional response circuit as described above and elsewhere herein. The system evaluation circuit 9408 may periodically share data with the communication circuit 9442 for transmittal to the remote server 9440 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 9446. Because relevant operating conditions and/or failure modes may occur as sensor values approach one or more criteria, the system evaluation circuit 9408 may share data with the communication circuit 9462 for transmittal to the remote server 9440 based on the fit of data relative to one or more criteria. Based on one sensor input meeting or approaching specified criteria or range, the system evaluation circuit 9408 may share additional data such as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The system evaluation circuit 9408 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server. In example embodiments, as shown in Figure 87, the communications circuit 9442 may communicate data directly to a remote server 9440. In example embodiments, as shown in Figure 88, the communications circuit 9442 may communicate data to an intermediate computer 9450 which may include a processor 9452 running an operating system 9454 and a data storage circuit 9456.

[0827] In example embodiments, as illustrated in Figures 89 and 90, a data collection system 9458 may have a plurality of monitoring devices 9448 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment (both the same and different types of equipment) in the same facility as well as collecting data from monitoring devices in multiple facilities. A monitoring application 9446 on a remote server 9440 may receive and store one or more of detection values, timing signals and data coming from a plurality of the various monitoring devices 9448. In example embodiments, as shown in Figure 89, the communications circuit 9442 may communicate data directly to a remote server 9440. In example embodiments, as shown in Figure 90, the communications circuit 9442 may communicate data to an intermediate computer 9450, which may include a processor 9452 running an operating system 9454 and a data storage circuit 9456. There may be an individual

intermediate computer 9450 associated with each monitoring device 9264 or an individual intermediate computer 9450 may be associated with a plurality of monitoring devices 9448 where the intermediate computer 9450 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 9440.

[0828] The monitoring application 9446 may select subsets of detection values, timing signals, data, product performance and the like to be jointly analyzed. Subsets for analysis may be selected based on component type, component materials, or a single type of equipment in which a component is operating. Subsets for analysis may be selected or grouped based on common operating conditions or operational history such as size of load, operational condition (e.g., intermittent, continuous), operating speed or tachometer, common ambient environmental conditions such as humidity, temperature, air or fluid particulate, and the like. Subsets for analysis may be selected based on common anticipated state information. Subsets for analysis may be selected based on the effects of other nearby equipment such as nearby machines rotating at similar frequencies, nearby equipment producing electromagnetic fields, nearby equipment producing heat, nearby equipment inducing movement or vibration, nearby equipment emitting vapors, chemicals or particulates, or other potentially interfering or intervening effects.

[0829] The monitoring application 9446 may analyze a selected subset. In an illustrative example, data from a single component may be analyzed over different time periods such as one operating cycle, cycle to cycle comparisons, trends over several operating cycles/time such as a month, a year, the life of the component or the like. Data from multiple components of the same type may also be analyzed over different time periods. Trends in the data such as changes in frequency or amplitude may be correlated with failure and maintenance records associated with the same component or piece of equipment. Trends in the data such as changing rates of change associated with start-up or different points in the process may be identified. Additional data may be introduced into the analysis such as output product quality, output quantity (such as per unit of time), indicated success or failure of a process, and the like. Correlation of trends and values for different types of data may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected performance. The analysis may identify model improvements to the model for anticipated state information, recommendations around sensors to be used, positioning of sensors and the like. The analysis may identify additional data to collect and store. The analysis may identify recommendations regarding needed maintenance and repair and/or the scheduling of preventative maintenance. The analysis may identify recommendations around purchasing replacement components and the timing of the replacement of the components. The analysis may identify recommendations regarding future geometry changes to reduce torsion on components. The analysis may result in warning regarding dangers of catastrophic failure conditions. This information may be transmitted back to the monitoring device to update types of data collected and analyzed locally or to influence the design of future monitoring devices.

[0830] In example embodiments, the monitoring application 9446 may have access to equipment specifications, equipment geometry, component specifications, component materials, anticipated state information for a plurality of component types, operational history, historical detection

values, component life models and the like for use analyzing the selected subset using rule-based or model-based analysis. In example embodiments, the monitoring application 9446 may feed a neural net with the selected subset to learn to recognize various operating states, health states (e.g., lifetime predictions) and fault states utilizing deep learning techniques. In example embodiments, a hybrid of the two techniques (model-based learning and deep learning) may be used.

[0831] In an illustrative and non-limiting example, the health of the rotating components on conveyors and lifters in an assembly line may be monitored using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0832] In an illustrative and non-limiting example, the health of the rotating components in water pumps on industrial vehicles may be monitored using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0833] In an illustrative and non-limiting example, the health of rotating components in compressors in gas handling systems may be monitored using the data monitoring devices and data collection systems described herein.

[0834] In an illustrative and non-limiting example, the health of the rotating components in compressors situated in the gas and oil fields may be monitored using the data monitoring devices and data collection systems described herein.

[0835] In an illustrative and non-limiting example, the health of the rotating components in factory air conditioning units may be evaluated using the techniques, data monitoring devices and data collection systems described herein.

[0836] In an illustrative and non-limiting example, the health of the rotating components in factory mineral pumps may be evaluated using the techniques, data monitoring devices and data collection systems described herein.

[0837] In an illustrative and non-limiting example, the health of the rotating components such as shafts, bearings, and gears in drilling machines and screw drivers situated in the oil and gas fields may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0838] In an illustrative and non-limiting example, the health of rotating components such as shafts, bearings, gears, and rotors of motors situated in the oil and gas fields may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0839] In an illustrative and non-limiting example, the health of rotating components such as blades, screws and other components of pumps situated in the oil and gas fields may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0840] In an illustrative and non-limiting example, the health of rotating components such as shafts, bearings, motors, rotors, stators, gears, and other components of vibrating conveyors situated in the oil and gas fields may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0841] In an illustrative and non-limiting example, the health of rotating components such as bearings, shafts, motors, rotors, stators, gears, and other components of mixers situated in the oil and gas fields may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0842] In an illustrative and non-limiting example, the health of rotating components such as bearings, shafts, motors, rotors, stators, gears, and other components of centrifuges situated in oil and gas refineries may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0843] In an illustrative and non-limiting example, the health of rotating components such as bearings, shafts, motors, rotors, stators, gears, and other components of refining tanks situated in oil and gas refineries may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0844] In an illustrative and non-limiting example, the health of rotating components such as bearings, shafts, motors, rotors, stators, gears, and other components of rotating tank/mixer agitators to promote chemical reactions deployed in chemical and pharmaceutical production lines may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0845] In an illustrative and non-limiting example, the health of rotating components such as bearings, shafts, motors, rotors, stators, gears, and other components of mechanical/rotating agitators to promote chemical reactions deployed in chemical and pharmaceutical production lines may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0846] In an illustrative and non-limiting example, the health of rotating components such as bearings, shafts, motors, rotors, stators, gears, and other components of propeller agitators to promote chemical reactions deployed in chemical and pharmaceutical production lines may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0847] In an illustrative and non-limiting example, the health of bearings and associated shafts, motors, rotors, stators, gears, and other components of vehicle steering mechanisms may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0848] In an illustrative and non-limiting example, the health of bearings and associated shafts, motors, rotors, stators, gears, and other components of vehicle engines may be evaluated using the torsional analysis techniques, data monitoring devices and data collection systems described herein.

[0849] In example embodiments, a monitoring device for estimating an anticipated lifetime of a rotating component in an industrial machine may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor

positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in the identification of torsional vibration; and a system analysis circuit structured to utilize the identified torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify an anticipated lifetime of the rotating component. In example embodiments, the monitoring device may further comprise a response circuit to perform at least one operation in response to the anticipated lifetime of the rotating component, wherein the plurality of input sensors includes at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor, a tachometer, and the like. At least one operation may comprise issuing at least one of an alert and a warning, storing additional data in the data storage circuit, ordering a replacement of the rotating component, scheduling replacement of the rotating component, recommending alternatives to the rotating component, and the like.

[0850] In example embodiments, a monitoring device for evaluating the health of a rotating component in an industrial machine may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in the identification of torsional vibration; and a system analysis circuit structured to utilize the identified torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify the health of the rotating component. In example embodiments, the monitoring device may further comprise a response circuit to perform at least one operation in response to the health of the rotating component. The plurality of input sensors may include at least two sensors selected from the group consisting of a temperature sensor, a load sensor, an optical vibration sensor, an acoustic wave sensor, a heat flux sensor, an infrared sensor, an accelerometer, a tri-axial vibration sensor a tachometer, and the like. The monitoring device may issue an alert and an alarm, such as the at least one operation storing additional data in the data storage circuit, ordering a replacement of the rotating component, scheduling replacement of the rotating component, recommending alternatives to the rotating component, and the like.

[0851] In example embodiments, a monitoring device for evaluating the operational state of a rotating component in an industrial machine may comprise a data acquisition circuit structured to

interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in the identification of torsional vibration; and a system analysis circuit structured to utilize the identified torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify the operational state of the rotating component. In example embodiments, the operational state may be a current or future operational state. A response circuit may perform at least one operation in response to the operational state of the rotating component. The at least one operation may store additional data in the data storage circuit, order a replacement of the rotating component, schedule a replacement of the rotating component, recommending alternatives to the rotating component, and the like.

[0852] In example embodiments, s monitoring device for evaluating the operational state of a rotating component in an industrial machine may include a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in the identification of torsional vibration; and a system analysis circuit structured to utilize the identified torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify the operational state of the rotating component, wherein the data acquisition circuit comprises a multiplexer circuit whereby alternative combinations of the detection values may be selected based on at least one of user input, a detected state and a selected operating parameter for a machine. The operational state may be a current or future operational state. The at least one operation may enable or disable one or more portions of the multiplexer circuit, or alter the multiplexer control lines. The data acquisition circuit may include at least two multiplexer circuits and the at least one operation comprises changing connections between the at least two multiplexer circuits.

[0853] In example embodiments, a system for evaluating an operational state a rotating component in a piece of equipment may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular

position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in identification of any torsional vibration; a system analysis circuit structured to utilize the torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify the operational state of the rotating component; and a communication module enabled to communicate the operational state of the rotating component, the torsional vibration and detection values to a remote server, wherein the detection values communicated are based partly on the operational state of the rotating component and the torsional vibration; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the monitoring devices. The analysis of the subset of detection values may include transitory signal analysis to identify the presence of high frequency torsional vibration. The monitoring application may be structured to subset detection values based on one of: operational state, torsional vibration, type of the rotating component, operational conditions under which detection values were measured, and type or equipment. The analysis of the subset of detection values may include feeding a neural net with the subset of detection values and supplemental information to learn to recognize various operating states, health states and fault states utilizing deep learning techniques. The supplemental information may include one of component specification, component performance, equipment specification, equipment performance, maintenance records, repair records an anticipated state model, and the like. The operational state may include a current or future operational state. The monitoring device may include a response circuit to perform at least one operation in response to the operational state of the rotating component. The at least one operation may include storing additional data in the data storage circuit.

[0854] In example embodiments, a system for evaluating the health of a rotating component in a piece of equipment may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of: an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in identification of torsional vibration; a system analysis circuit structured to utilize the torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify the health of the rotating component; and a communication module enabled to

communicate the health of the rotating component, the torsional vibrations and detection values to a remote server, wherein the detection values communicated are based partly on the health of the rotating component and the torsional vibration; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the monitoring devices. In example embodiments, the analysis of the subset of detection values may include transitory signal analysis to identify the presence of high frequency torsional vibration. The monitoring application may be structured to subset detection values. The analysis of the subset of detection values may include feeding a neural net with the subset of detection values and supplemental information to learn to recognize various operating states, health states and fault states utilizing deep learning techniques. The supplemental information may include one of component specification, component performance, equipment specification, equipment performance, maintenance records, repair records and an anticipated state model. The operational state may be a current or future operational state. A response circuit may perform at least one operation in response to the health of the rotating component.

[0855] In example embodiments, a system for estimating an anticipated lifetime of a rotating component in a piece of equipment may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in identification of torsional vibration; a system analysis circuit structured to utilize the torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify an anticipated life the rotating component; and a communication module enabled to communicate the anticipated life of the rotating component, the torsional vibrations and detection values to a remote server, wherein the detection values communicated are based partly on the anticipated life of the rotating component and the torsional vibration; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the monitoring devices. In example embodiments, the analysis of the subset of detection values may include transitory signal analysis to identify the presence of high frequency torsional vibration. The monitoring application may be structured to subset detection values based on one of anticipated life of the rotating component, torsional vibration, type of the rotating component, operational conditions under which detection values were measured, and type of equipment. The analysis of the subset of detection values may include feeding a neural net with the subset of detection values and supplemental information to learn to recognize various operating states, health states, life expectancies and fault states utilizing deep learning techniques. The supplemental information may include one of component specification, component performance,

equipment specification, equipment performance, maintenance records, repair records and an anticipated state model. The monitoring device may include a response circuit to perform at least one operation in response to the anticipated life of the rotating component. The at least one operation may include one of ordering a replacement of the rotating component, scheduling replacement of the rotating component, and recommending alternatives to the rotating component.

[0856] In example embodiments, a system for evaluating the health of a variable frequency motor in an industrial environment may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a data storage circuit structured to store specifications, system geometry, and anticipated state information for a plurality of rotating components, store historical component performance and buffer the plurality of detection values for a predetermined length of time; and a torsional analysis circuit structured to utilize transitory signal analysis to analyze the buffered detection values relative to the rotating component specifications and anticipated state information resulting in identification of torsional vibration; a system analysis circuit structured to utilize the torsional vibration and at least one of an anticipated state, historical data and a system geometry to identify a motor health parameter; and a communication module enabled to communicate the motor health parameter, the torsional vibrations and detection values to a remote server, wherein the detection values communicated are based partly on the motor health parameter and the torsional vibration; and a monitoring application on the remote server structured to receive, store and jointly analyze a subset of the detection values from the monitoring devices.

[0857] In example embodiments, a system for data collection, processing, and torsional analysis of a rotating component in an industrial environment may comprise a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of a plurality of input sensors, wherein the plurality of input sensors comprises at least one of an angular position sensor, an angular velocity sensor and an angular acceleration sensor positioned to measure the rotating component; a streaming circuit for streaming at least a subset of the acquired detection values to a remote learning system; and a remote learning system including a torsional analysis circuit structured to analyze the detection values relative to a machine-based understanding of the state of the at least one rotating component. The machine-based understanding may be developed based on a model of the rotating component that determines a state of the at least one rotating component based at least in part on the relationship of the behavior of the rotating component to an operating frequency of a component of the industrial machine. The state of the at least one rotating component may be at least one of an operating state, a health state, a predicted lifetime state and a fault state. The machine-based understanding may be developed based by providing inputs to a deep learning machine, wherein the inputs comprise a plurality of streams of detection values for a plurality of rotating components and a plurality of measured state values for the plurality of rotating components. The state of the at least one rotating component may be at least one of an operating state, a health state, a predicted lifetime state and a fault state.

[0858] In example embodiments, information about the health or other status or state information of or regarding a component or piece of industrial equipment may be obtained by monitoring the condition of various components throughout a process. Monitoring may include monitoring the amplitude of a sensor signal measuring attributes such as temperature, humidity, acceleration, displacement and the like. An embodiment of a data monitoring device 9700 is shown in Figure 91 and may include a plurality of sensors 9706 communicatively coupled to a controller 9702. The controller 9702 may include a data acquisition circuit 9704, a signal evaluation circuit 9708, a data storage circuit 9716 and a response circuit 9710. The signal evaluation circuit 9708 may comprise a circuit for detecting a fault in one or more sensors, or a set of sensors, such as an overload detection circuit 9712, a sensor fault detection circuit 9714, or both. Additionally, the signal evaluation circuit 9708 may optionally comprise one or more of a peak detection circuit, a phase detection circuit, a bandpass filter circuit, a frequency transformation circuit, a frequency analysis circuit, a phase lock loop circuit, a torsional analysis circuit, a bearing analysis circuit, and the like.

[0859] The plurality of sensors 9706 may be wired to ports on the data acquisition circuit 9704. The plurality of sensors 9706 may be wirelessly connected to the data acquisition circuit 9704. The data acquisition circuit 9704 may be able to access detection values corresponding to the output of at least one of the plurality of sensors 9706 where the sensors 9706 may be capturing data on different operational aspects of a piece of equipment or an operating component.

[0860] The selection of the plurality of sensors 9706 for a data monitoring device 9700 designed for a specific component or piece of equipment may depend on a variety of considerations such as accessibility for installing new sensors, incorporation of sensors in the initial design, anticipated operational and failure conditions, resolution desired at various positions in a process or plant, reliability of the sensors, and the like. The impact of a failure, time response of a failure (e.g., warning time and/or off-nominal modes occurring before failure), likelihood of failure, and/or sensitivity required and/or difficulty to detection failure conditions may drive the extent to which a component or piece of equipment is monitored with more sensors and/or higher capability sensors being dedicated to systems where unexpected or undetected failure would be costly or have severe consequences.

[0861] Depending on the type of equipment, the component being measured, the environment in which the equipment is operating and the like, sensors 9706 may comprise, without limitation, one or more of the following: a vibration sensor, a thermometer, a hygrometer, a voltage sensor and/or a current sensor (for the component and/or other sensors measuring the component), an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, a thermal imager, an acoustic wave sensor, a displacement sensor, a turbidity meter, a viscosity meter, a axial load sensor, a radial load sensor, a tri-axial sensor, an accelerometer, a speedometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an optical (laser) particle counter, an ultrasonic sensor, an acoustical sensor, a heat flux sensor, a galvanic sensor, a

magnetometer, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference.

[0862] The sensors 9706 may provide a stream of data over time that has a phase component, such as relating to acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece of equipment or an operating component. The sensors 9706 may provide a stream of data that is not conventionally phase-based, such as temperature, humidity, load, and the like. The sensors 9706 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[0863] In example embodiments, as illustrated in Figure 91, the sensors 9706 may be part of the data monitoring device 9700, referred to herein in some cases as a data collector, which in some cases may comprise a mobile or portable data collector. In example embodiments, as illustrated in Figures 92 and 93, one or more external sensors 9724, which are not explicitly part of a monitoring device 9718 but rather are new, previously attached to or integrated into the equipment or component, may be opportunistically connected to or accessed by the monitoring device 9718. The monitoring device may include a data acquisition circuit 9722, a signal evaluation circuit 9708, a data storage circuit 9716 and a response circuit 9710. The signal evaluation circuit 9708 may comprise an overload detection circuit 9712, a sensor fault detection circuit 9714, or both. Additionally, the signal evaluation circuit 9708 may optionally comprise one or more of a peak detection circuit, a phase detection circuit, a bandpass filter circuit, a frequency transformation circuit, a frequency analysis circuit, a phase lock loop circuit, a torsional analysis circuit, a bearing analysis circuit, and the like. The data acquisition circuit 9722 may include one or more input ports 9726.

[0864] The one or more external sensors 9724 may be directly connected to the one or more input ports 9726 on the data acquisition circuit 9722 of the controller 9720 or may be accessed by the data acquisition circuit 9722 wirelessly, such as by a reader, interrogator, or other wireless connection, such as over a short-distance wireless protocol. In example embodiments, as shown in Figure 93, a data acquisition circuit 9722 may further comprise a wireless communication circuit 9730. The data acquisition circuit 9722 may use the wireless communication circuit 9730 to access detection values corresponding to the one or more external sensors 9724 wirelessly or via a separate source or some combination of these methods.

[0865] In example embodiments, the data storage circuit 9716 may be structured to store sensor specifications, anticipated state information and detected values. The data storage circuit 9716 may provide specifications and anticipated state information to the signal evaluation circuit 9708.

[0866] In example embodiments, an overload detection circuit 9712 may detect sensor overload by comparing the detected value associated with the sensor with a detected value associated with a sensor having a greater range/lower resolution monitoring the same component/attribute. Inconsistencies in measured value may indicate that the higher resolution sensor may be overloaded. In example embodiments, an overload detection circuit 9712 may detect sensor overload by evaluating consistency of sensor reading with readings from other sensor data

(monitoring the same or different aspects of the component/piece of equipment. In example embodiments, an overload detection circuit 9712 may detect sensor overload by evaluating data collected by other sensors to identify conditions likely to result in sensor overload (e.g., heat flux sensor data indicative of the likelihood of overloading a sensor in a given location, accelerometer data indicating a likelihood of overloading a velocity sensor, and the like). In example embodiments, an overload detection circuit 9712 may detect sensor overload by identifying flat line output following a rising trend. In example embodiments, an overload detection circuit 9712 may detect sensor overload by transforming the sensor data to frequency data, using for example a Fast Fourier Transform (FFT), and then looking for a "ski-jump" in the frequency data which may result from the data being clipped due to an overloaded sensor. A sensor fault detection circuit 9714 may identify failure of the sensor itself, sensor health, or potential concerns regarding validity of sensor data. Rate of value change may be used to identify failure of the sensor itself. For example, a sudden jump to a maximum output may indicate a failure in the sensor rather than an overload of the sensor. In example embodiments, an overload detection circuit 9712 and/or a sensor fault detection circuit 9712 may utilize sensor specifications, anticipated state information, sensor models and the like in the identification of sensor overload, failure, error, invalid data, and the like. In example embodiments, the overload detection circuit 9712 or the sensor fault detection circuit 9714 may use detection values from other sensors and output from additional components such as a peak detection circuit and/or a phase detection circuit and/or a bandpass filter circuit and/or a frequency transformation circuit and/or a frequency analysis circuit and/or a phase lock loop circuit and the like to identify potential sources for the identified sensor overload, sensor faults, sensor failure, or the like. Sources or factors involved in sensor overload may include limitations on sensor range, sensor resolution, and sensor sampling frequency. Sources of apparent sensor overload may be due to a range, resolution or sampling frequency of a multiplexor supplying detection values associated with the sensor. Sources of factors involved in apparent sensor faults or failures may include environmental conditions; for example, excessive heat or cold may be associated with damage to semiconductor-based sensors, which may result in erratic sensor data, failure of a sensor to produce data, data that appears out of the range of normal behavior (e.g., large, discrete jumps in temperature for a system that does not normally experience such changes). Surges in current and/or voltage may be associated with damage to electrically connected sensors with sensitive components. Excessive vibration may result in physical damage to sensitive components of a sensor such as wires and/or connectors. An impact, which may be indicated by sudden acceleration or acoustical data may result in physical damage to a sensor with sensitive components such as wires and/or connectors. A rapid increase in humidity in the environment surrounding a sensor or an absence of oxygen may indicate water damage to a sensor. A sudden absence of signal from a sensor may be indicative of sensor disconnection which may be due to vibration, impact and the like. A sensor that requires power may run out of battery power or be disconnected from a power source. In example embodiments, the overload detection circuit 9712 or the sensor fault detection circuit 9714 may output a sensor status where the sensor status may be one of sensor overload,

sensor failure, sensor fault, sensor healthy, and the like. The sensor fault detection circuit 9714 may determine one of a sensor fault status and a sensor validity status.

[0867] In example embodiments, as illustrated in Figure 94, the data acquisition circuit 9722 may further comprise a multiplexer circuit 9731 as described elsewhere herein. Outputs from the multiplexer circuit 9731 may be utilized by the signal evaluation circuit 9708. The response circuit 9710 may have the ability to turn on or off portions of the multiplexor circuit 9731. The response circuit 9710 may have the ability to control the control channels of the multiplexor circuit 9731.

[0868] In example embodiments, the response circuit 9710 may initiate a variety of actions based on the sensor status provided by the overload detection circuit 9712. The response circuit 9710 may continue using the sensor if the sensor status is "sensor healthy." The response circuit 9710 may adjust a sensor scaling value (e.g., from 100mV/gram to 10 mV/gram). The response circuit 9710 may increase an acquisition range for an alternate sensor. The response circuit 9710 may back sensor data out of previous calculations and evaluations such as bearing analysis, torsional analysis and the like. The response circuit 9710 may use projected or anticipated data (based on data acquired prior to overload/failure) in place of the actual sensor data for calculations and evaluations such as bearing analysis, torsional analysis and the like. The response circuit 9710 may issue an alarm. The response circuit 9710 may issue an alert that may comprise notification that the sensor is out of range together with information regarding the extent of the overload such as "overload range- data response may not be reliable and/or linear", "destructive range- sensor may be damaged," and the like. The response circuit 9710 may issue an alert where the alert may comprise information regarding the effect of sensor load such as "unable to monitor machine health" due to sensor overload/failure," and the like.

[0869] In example embodiments, the response circuit 9710 may cause the data acquisition circuit 9704 to enable or disable the processing of detection values corresponding to certain sensors based on the sensor statues described above. This may include switching to sensors having different response rates, sensitivity, ranges, and the like; accessing new sensors or types of sensors, accessing data from multiple sensors, recruiting additional data collectors (such as routing the collectors to a point of work, using routing methods and systems disclosed throughout this disclosure and the documents incorporated by reference) and the like. Switching may be undertaken based on a model, a set of rules, or the like. In example embodiments, switching may be under control of a machine learning system, such that switching is controlled based on one or more metrics of success, combined with input data, over a set of trials, which may occur under supervision of a human supervisor or under control of an automated system. Switching may involve switching from one input port to another (such as to switch from one sensor to another). Switching may involve altering the multiplexing of data, such as combining different streams under different circumstances. Switching may involve activating a system to obtain additional data, such as moving a mobile system (such as a robotic or drone system), to a location where different or additional data is available (such as positioning an image sensor for a different view or positioning a sonar sensor for a different direction of collection) or to a location where different sensors can be accessed (such as moving a collector to connect up to a sensor that is disposed at a location in

an environment by a wired or wireless connection). This switching may be implemented by changing the control signals for a multiplexor circuit 9731 and/or by turning on or off certain input sections of the multiplexor circuit 9731.

[0870] In example embodiments, the response circuit 9710 may make recommendations for the replacement of certain sensors in the future with sensors having different response rates, sensitivity, ranges, and the like. The response circuit 9710 may recommend design alterations for future embodiments of the component, the piece of equipment, the operating conditions, the process, and the like.

[0871] In example embodiments, the response circuit 9710 may recommend maintenance at an upcoming process stop or initiate a maintenance call where the maintenance may include the replacement of the sensor with the same or an alternate type of sensor having a different response rate, sensitivity, range and the like. In example embodiments, the response circuit 9710 may implement or recommend process changes – for example to lower the utilization of a component that is near a maintenance interval, operating off-nominally, or failed for purpose but still at least partially operational, to change the operating speed of a component (such as to put it in a lower-demand mode), to initiate amelioration of an issue (such as to signal for additional lubrication of a roller bearing set, or to signal for an alignment process for a system that is out of balance), and the like.

[0872] In example embodiments, the signal evaluation circuit 9708 and/or the response circuit 9710 may periodically store certain detection values in the data storage circuit 9716 to enable the tracking of component performance over time. In example embodiments, based on sensor status, as described elsewhere herein recently measured sensor data and related operating conditions such as RPMs, component loads, temperatures, pressures, vibrations or other sensor data of the types described throughout this disclosure in the data storage circuit 9716 to enable the backing out of overloaded/failed sensor data. The signal evaluation circuit 9708 may store data at a higher data rate for greater granularity in future processing, the ability to reprocess at different sampling rates, and/or to enable diagnosing or post-processing of system information where operational data of interest is flagged, and the like.

[0873] In embodiments as shown in Figures 95, 96, 97, and 98, a data monitoring system 9726 may include at least one data monitoring device 9728. At least one data monitoring device 9728 may include sensors 9706 and a controller 9730 comprising a data acquisition circuit 9704, a signal evaluation circuit 9708, a data storage circuit 9716, and a communication circuit 9754 to allow data and analysis to be transmitted to a monitoring application 9736 on a remote server 9734. The signal evaluation circuit 9708 may include at least an overload detection circuit 9712. The signal evaluation circuit 9708 may periodically share data with the communication circuit 9732 for transmittal to the remote server 9734 to enable the tracking of component and equipment performance over time and under varying conditions by a monitoring application 9736. Based on the sensor status, the signal evaluation circuit 9708 and/or response circuit 9710 may share data with the communication circuit 9732 for transmittal to the remote server 9734 based on the fit of data relative to one or more criteria. Data may include recent sensor data and additional data such

as RPMs, component loads, temperatures, pressures, vibrations, and the like for transmittal. The signal evaluation circuit 9708 may share data at a higher data rate for transmittal to enable greater granularity in processing on the remote server.

[0874] In example embodiments, as shown in Figure 95, the communication circuit 9732 may communicate data directly to a remote server 9734. In embodiments as shown in Figure 96, the communication circuit 9732 may communicate data to an intermediate computer 9738 which may include a processor 9740 running an operating system 9742 and a data storage circuit 9744.

[0875] In example embodiments, as illustrated in Figures 97 and 98, a data collection system 9746 may have a plurality of monitoring devices 9728 collecting data on multiple components in a single piece of equipment, collecting data on the same component across a plurality of pieces of equipment, (both the same and different types of equipment) in the same facility as well as collecting data from monitoring devices in multiple facilities. A monitoring application 9736 on a remote server 9734 may receive and store one or more of detection values, timing signals and data coming from a plurality of the various monitoring devices 9728.

[0876] In example embodiments, as shown in Figure 97, the communication circuit 9732 may communicate data directly to a remote server 9734. In example embodiments, as shown in Figure 98, the communication circuit 9732 may communicate data to an intermediate computer 9738 which may include a processor 9740 running an operating system 9742 and a data storage circuit 9744. There may be an individual intermediate computer 9738 associated with each monitoring device 9728 or an individual intermediate computer 9738 may be associated with a plurality of monitoring devices 9728 where the intermediate computer 9738 may collect data from a plurality of data monitoring devices and send the cumulative data to the remote server 9734. Communication to the remote server 9734 may be streaming, batch (e.g., when a connection is available) or opportunistic.

[0877] The monitoring application 9736 may select subsets of the detection values to be jointly analyzed. Subsets for analysis may be selected based on a single type of sensor, component or a single type of equipment in which a component is operating. Subsets for analysis may be selected or grouped based on common operating conditions such as size of load, operational condition (e.g., intermittent, continuous), operating speed or tachometer, common ambient environmental conditions such as humidity, temperature, air or fluid particulate, and the like. Subsets for analysis may be selected based on the effects of other nearby equipment such as nearby machines rotating at similar frequencies, nearby equipment producing electromagnetic fields, nearby equipment producing heat, nearby equipment inducing movement or vibration, nearby equipment emitting vapors, chemicals or particulates, or other potentially interfering or intervening effects.

[0878] In example embodiments, the monitoring application 9736 may analyze the selected subset. In an illustrative example, data from a single sensor may be analyzed over different time periods such as one operating cycle, several operating cycles, a month, a year, the life of the component or the like. Data from multiple sensors of a common type measuring a common component type may also be analyzed over different time periods. Trends in the data such as changing rates of change associated with start-up or different points in the process may be identified. Correlation of trends

and values for different sensors may be analyzed to identify those parameters whose short-term analysis might provide the best prediction regarding expected sensor performance. This information may be transmitted back to the monitoring device to update sensor models, sensor selection, sensor range, sensor scaling, sensor sampling frequency, types of data collected and analyzed locally or to influence the design of future monitoring devices.

[0879] In example embodiments, the monitoring application 9736 may have access to equipment specifications, equipment geometry, component specifications, component materials, anticipated state information for a plurality of sensors, operational history, historical detection values, sensor life models and the like for use analyzing the selected subset using rule-based or model-based analysis. The monitoring application 9736 may provide recommendations regarding sensor selection, additional data to collect, or data to store with sensor data. The monitoring application 9736 may provide recommendations regarding scheduling repairs and/or maintenance. The monitoring application 9736 may provide recommendations regarding replacing a sensor. The replacement sensor may match the sensor being replaced or the replacement sensor may have a different range, sensitivity, sampling frequency and the like.

[0880] In example embodiments, the monitoring application 9736 may include a remote learning circuit structured to analyze sensor status data (e.g., sensor overload, sensor faults, sensor failure) together with data from other sensors, failure data on components being monitored, equipment being monitored, product being produced, and the like. The remote learning system may identify correlations between sensor overload and data from other sensors.

[0881] Within the data collection, monitoring, and control environment of the industrial IoT are large and various sensor sets, which make efficient setup and timely changes to sensor data collection a challenge. Continuous collection from all sensors may be impossible given the large number of sensors and limited resources, such as limited availability of power and limited data collection and management facilities, including various limitations in availability and performance of sensor data collection devices, input/output interfaces, data transfer facilities, data storage, data analysis facilities, and the like. The number of sensors collected from at any given time must therefore be limited in an intelligent but timely manner, both at the time of setting up initial collection and during the process of collection, including handling rapid changes to a present collection scheme based on a change in state of a system, operational conditions (e.g., an alert condition, change in operational mode, etc.), or the like. Embodiments of the methods and systems disclosed herein may therefore include rapid route creation and modification for routing collectors, such as by taking advantage of hierarchical templates, execution of smart route changes, monitoring and responding to changes in operational conditions, and the like.

[0882] In example embodiments, rapid route creation and modification for data collection in an industrial environment may take advantage of hierarchical templates. Templates may be used to take advantage of 'like' machinery that can utilize the same hierarchical sensor routing scheme. For example, among many possible types of machines about which data may be collected, the members of a certain class of motor, such as a stepper motor class, may have very similar sensor routing needs, such as for routine operations, routine maintenance, and failure mode detection, that

may be described in a common hierarchy of sensor collection routines. The user installing a new stepper motor may then use the 'stepper motor hierarchical routing template' for the new motor. After installation, the stepper motor hierarchical routing template may then be used to change the routing schemes for changing conditions. The user may optionally make adjustments to the template as needed per unique motor functions, applications, environments, modes, and the like. The use of a template for deploying a routing scheme greatly reduces the time a user requires to configure the routing scheme for a new motor, or to deploy new routing technologies on an existing system that utilizes traditional sensor collection methods. Once the hierarchical routing template is in place, the sensor collection routine may be changed quickly based on the template, thus allowing for rapid route modification under changing conditions, such as: a change in the operating mode of the stepper motor that requires a different subset of sensors for monitoring, a limit alert or failure indication that requires a more focused subset of sensors for use in diagnosing the problem, and the like. Hierarchical routing templates thus allow for rapid deployment of sensor routing configurations, as well as allowing the sensed industrial environment to be altered dynamically as conditions change.

[0883] A functional hierarchy of routing templates may include different hierarchical configurations for a component, machine, system, industrial environment, and the like, including all sensors and a plurality of configurations formed from a subset of all sensors. At a system level, an 'all-sensor' configuration may include: a connection map to all sensors in a system, mapping to all onboard instrumentation sensors (e.g., monitoring points reporting within a machine or set of machines), mapping to an environment's sensors (e.g., monitoring points around the machines/equipment, but not necessarily onboard), mapping to available sensors on data collectors (e.g., data collectors that can be flexibly provisioned for particular data among different kinds), a unified map combining different individual mappings, and the like. A routing configuration may be provided, such as to indicate how to implement an operational routing scheme, a scheduled maintenance routing scheme (e.g., collecting from a greater set of overall sensors than in operational mode, but distributed across the system, or a focused sensor set for specific components, functions, and modes), one or more failure mode routing schemes for multiple focused sensor collection groups targeting different failure mode analyses (e.g., for a motor, one failure mode may be for bearings, another for startup speed-torque, where a different subset of sensor data is needed based on the failure mode, such as detected in anomalous readings taken during operations or maintenance), power savings (e.g., weather conditions necessitating reduced plant power), and the like.

[0884] As noted, hierarchical templates may also be conditional (e.g., rule-based), such as templates with conditional routing based on parameters, such as sensed data during a first collection period, where a subsequent routing configuration is varied. Within the hierarchy, nodes in a graph or tree may indicate forks by which conditional logic may be used, such as to select a given subset of sensors for a given operational mode. Thus, the hierarchical template may be associated with a rule-based or model-based expert system, which may facilitate automated routing based on the hierarchical template and based on observed conditions, such as based on a type of

machine and its operational state, environmental context, or the like. In a non-limiting example, a hierarchical template may have an initial collection configuration and a conditional hierarchy in place to switch from the initial collection configuration to a second collection configuration based on the sensed conditions of an initial sensor collection. Continuing this example, among various possible machines, a conveyor system may have a plurality of sensors for collection in an initial collection, but once the first data is collected and analyzed, if the conveyor is determined to be in an idle state (such as due to the absence of a signal above a minimum threshold on a motion sensor), then the system may switch to a sensor data collection regime that is appropriate for the idles state of the conveyor (e.g., using a very small subset of the plurality of sensors, such as just using the motion sensor to detect departure from the idle state, at which point the original regime may be renewed and the rest of a sensor set may be re-engaged). Thus, when the collection of sensor data detects a changed condition to a state, an operational mode, an environmental condition, or the like, the sensor data collection may be switched to an appropriate configuration.

[0885] Hierarchical templates for one collector may be based on coordination of routing with that of other collectors. For instance, a collector might be set up to perform vibration analysis while another collector is set up to perform pressure or temperature on each machine in a set of similar machines, rather than having each machine collect all of the data on each machine, where otherwise setup for different sensor types may be required for each collector for each machine. Factors such as the duration of sampling required, the time required to set up a given sensor, the amount of power consumed, the time available for collection as a whole, the data rate of input/output of a sensor and/or the collector, the bandwidth of a channel (wired or wireless) available for transmission of collected data, and the like can be considered in arranging the coordination of the routing of two or more collectors, such that various parallel and serial configurations may be undertaken to achieve an overall effectiveness. This may include optimizing the coordination using an expert system, such as a rule-based optimization, a model-based optimization, or optimization using machine learning.

[0886] A machine learning system may create a hierarchical template structure for improved routing, such as for teaching the system the default operating conditions (e.g., normal operations mode, systems online and average production), peak operations mode (max capability), slack production, and the like. The machine learning system may create a new hierarchical template based on monitored conditions, such as a template based on a production level profile, a rate of production profile, a detected failure mode pattern analysis, and the like. The application of a new machine learning created template may be based on a mode matching between current production conditions and a machine learning template condition (e.g., the machine learning system creates a new template for a new production profile, and applies that new template whenever that new profile is detected).

[0887] Rapid route creation may be enabled using one or more hierarchical routing templates, such as when a routing template pre-establishes a routing scheme for different conditions, and when a trigger event executes a change in the sensor routing scheme to accommodate the condition. In example embodiments, the trigger event may be an automatic change in routing based on a trigger

that indicates a possible failure mode that forces a change in routing scheme from operational to failure mode analysis; a human-executed change in routing scheme based on received sensor data; a learned routing change based on machine learning of when to trigger a change (e.g., as based on a machine being fed with a set of human-executed or human-supervised changes); a manual routing change (e.g., optional to automatic/rapid automatic change); a human executed change based on observed device performance; and the like. Routing changes may include: changing from an operational mode to an accelerated maintenance, a failure mode analysis, a power saving mode a high-performance/high-output mode (e.g., for peak power in a generation plant), and the like.

[0888] Switching hierarchical template configurations may be executed based on connectivity with end-device sensors. In a highly automated collection routing environment (e.g., an indoor networked assembly plant) different routing collection configurations may be employed for fixed and flexible industrial layouts. In a fixed industrial layout, such as a layout with a high degree of wired connectivity between end-device sensors, automated collectors, and networks, there may be different routing configurations for a network routing hierarchy portion, a collector sensor-collection hierarchy portion, a storage portion, and the like. For a more flexible industrial layout with various wired and wireless connections between end-device sensors, automated collectors, and networks, there may be different schemes. For instance, a moderately automated collection routing environment may include: automatic collection and periodic network connection; a robot-carried collector for periodic collection (e.g., a ground-based robot, a drone, an underwater device, a robot with network connection, a robot with intermittent network connection, a robot that periodically uploads collection); a routing scheme with periodic collection and automated routing; a scheme only collecting periodically but routed directly upon collection; a routing scheme with periodic collection and periodic automated routing to collect periodically; and, over longer periods of time, periodically routing multiple collections; and the like. For a lower degree of automated collection routing, there may be a combination of: automatic collection and human-aided collectors (e.g., humans collecting alone, humans aided by robots), scheduled collection and human-aided collectors (e.g., humans initiating collection, humans aided by robots for collection initiation, human launching a drone to collect data at a remote site), and the like.

[0889] In example embodiments, and referring to Figure 99, hierarchical templates may be utilized by a local data collection system 10500 for collection and monitoring of data collected through a plurality of input channels 10500, such as data from sensors 10514, IoT devices 10516, and the like. The local collection system 10512, also referred to herein as a data collector 10512, may comprise a data storage 10502; a data acquisition circuit 10504; a data analysis circuit 10506; and the like, wherein the monitoring facilities may be deployed: locally on the data collector 10512; in part locally on the data collector and in part on a remote information technology infrastructure component apart from the data collector; and the like. A monitoring system may comprise a plurality of input channels communicatively coupled to the data collector 10512. The data storage 10502 may be structured to store a plurality of collector route templates 10510 and sensor specifications for sensors 10514 that correspond to the input channels 10500, wherein the plurality of collector route templates 10510 each comprise a different sensor collection routine. A data

acquisition circuit 10504 may be structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels 10500, and a data analysis circuit 10506 structured to receive output data from the plurality of input channels 10500 and evaluate a current routing template collection routine based on the received output data, wherein the data collector 10500 is configured to switch from the current routing template collection routine to an alternative routing template collection routine based on the content of the output data. The monitoring system may further utilize a machine learning system (e.g., a neural network expert system), rule-based templates (e.g., based on an operational state of a machine with respect to which the input channels provide information, the input channels provide information, the input channels provide information), smart route changes, alarm states, network connectivity, self-organization amongst a plurality of data collectors, coordination of sensor groups, and the like.

[0890] In example embodiments, evaluation of the current routing templates may be based on operational mode routing collection schemes, such as a normal operational mode, a peak operational mode, an idle operational mode, a maintenance operational mode, a power saving operational mode, and the like. As a result of monitoring, the data collector may switch from a current routing template collection routine because the data analysis circuit determines a change in operating modes, such as the operating mode changing from an operational mode to an accelerated maintenance mode, the operating mode changing from an operational mode to a failure mode analysis mode, the operating mode changing from an operational mode to a power-saving mode, the operating mode changing from an operational mode to a high-performance mode, and the like. The data collector may switch from a current routing template collection routine based on a sensed change in a mode of operation, such as a failure condition, a performance condition, a power condition, a temperature condition, a vibration condition, and the like. The evaluation of the current routing template collection routine may be based on a collection routine with respect to a collection parameter, such as network availability, sensor availability, a time-based collection routine (e.g., on a schedule, over time), and the like.

[0891] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collector communicatively coupled to a plurality of input channels; a data storage structured to store a plurality of collector route templates and sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and a data analysis circuit structured to receive output data from the plurality of input channels and evaluate a current routing template collection routine based on the received output data, wherein the data collector is configured to switch from the current routing template collection routine to an alternative routing template collection routine based on the content of the output data. In example embodiments, the system is deployed locally on the data collector, in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, and the like. Each of the input channels may correspond to a sensor located in the environment. The evaluation of the current

routing template may be based on operational mode routing collection schemes. The operational mode is at least one of a normal operational mode, a peak operational mode, an idle operational mode, a maintenance operational mode, and a power saving operational mode. The data collector may switch from the current routing template collection routine because the data analysis circuit determines a change in operating modes, such as where the operating mode changes from an operational mode to an accelerated maintenance mode, from an operational mode to a failure mode analysis mode, from an operational mode to a power saving mode, from an operational mode to high-performance mode, and the like. The data collector may switch from the current routing template collection routine based on a sensed change in a mode of operation, such as where the sensed change is a failure condition, a performance condition, a power condition, a temperature condition, a vibration condition, and the like. The evaluation of the current routing template collection routine may be based on a collection routine with respect to a collection parameter, such as where the parameter is network availability, sensor availability, a time-based collection routine (e.g., where a routine collects sensor data on a schedule, evaluates sensor data over time).

[0892] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a plurality of collector route templates and sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and providing a data analysis circuit structured to receive output data from the plurality of input channels and evaluate a current routing template collection routine based on the received output data, wherein the data collector is configured to switch from the current routing template collection routine to an alternative routing template collection routine based on the content of the output data. In example embodiments, the computer-implemented method is deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment.

[0893] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a plurality of collector route templates and sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and providing a data analysis circuit structured to receive output data from the plurality of input channels and evaluate a current routing template collection routine based on the received output data, wherein the data collector is configured to switch from the current routing template collection routine to an alternative routing

template collection routine based on the content of the output data. In example embodiments, the instructions may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment.

[0894] In example embodiments, a monitoring system for data collection in an industrial environment may comprise a data collector communicatively coupled to a plurality of input channels; a data storage structured to store a plurality of collector route templates, sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and a machine learning data analysis circuit structured to receive output data from the plurality of input channels and evaluate a current routing template collection routine based on the received output data received over time, wherein the machine learning data analysis circuit learns received output data patterns, wherein the data collector is configured to switch from the current routing template collection routine to an alternative routing template collection routine based on the learned received output data patterns. In example embodiments, the monitoring system may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment. The machine learning data analysis circuit may include a neural network expert system. The evaluation of the current routing template may be based on operational mode routing collection schemes. The operational mode may be at least one of a normal operational mode, a peak operational mode, an idle operational mode, a maintenance operational mode, and a power saving operational mode. The data collector may switch from the current routing template collection routine because the data analysis circuit determines a change in operating modes, such as where the operating mode changes from an operational mode to an accelerated maintenance mode, from an operational mode to a failure mode analysis mode, from an operational mode to a power saving mode, from an operational mode to high-performance mode, and the like. The data collector may switch from the current routing template collection routine based on a sensed change in a mode of operation, such as where the sensed change is a failure condition, a performance condition, a power condition, a temperature condition, a vibration condition, and the like. The evaluation of the current routing template collection routine may be based on a collection routine with respect to a collection parameter, such as where the parameter is network availability, a sensor availability, a time-based collection routine (collects sensor data on a schedule, evaluates sensor data over time).

[0895] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a plurality of collector route templates, sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different

sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and providing a machine learning data analysis circuit structured to receive output data from the plurality of input channels and evaluate a current routing template collection routine based on the received output data received over time, wherein the machine learning data analysis circuit learns received output data patterns, wherein the data collector is configured to switch from the current routing template collection routine to an alternative routing template collection routine based on the learned received output data patterns. In example embodiments, the method may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment.

[0896] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a plurality of collector route templates, sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and providing a machine learning data analysis circuit structured to receive output data from the plurality of input channels and evaluate a current routing template collection routine based on the received output data received over time, wherein the machine learning data analysis circuit learns received output data patterns, wherein the data collector is configured to switch from the current routing template collection routine to an alternative routing template collection routine based on the learned received output data patterns. In example embodiments, the instructions may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment.

[0897] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collector communicatively coupled to a plurality of input channels; a data storage structured to store a collector route template, sensor specifications for sensors that correspond to the input channels, wherein the collector route template comprises a sensor collection routine; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and a data analysis circuit structured to receive output data from the plurality of input channels and evaluate the received output data with respect to a rule, wherein the data collector is configured to modify the sensor collection routine based on the application of the rule to the received output data. In example embodiments, the system may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to

a sensor located in the environment. The rule may be based on an operational state of a machine with respect to which the input channels provide information, on an anticipated state of a machine with respect to which the input channels provide information, on a detected fault condition of a machine with respect to which the input channels provide information, and the like. The evaluation of the received output data may be based on operational mode routing collection schemes, where the operational mode is at least one of a normal operational mode, a peak operational mode, an idle operational mode, a maintenance operational mode, and a power saving operational mode. The data collector may modify the sensor collection routine because the data analysis circuit determines a change in operating modes, such as where the operating mode changes from an operational mode to an accelerated maintenance mode, from an operational mode to a failure mode analysis mode, from an operational mode to a power saving mode, from an operational mode to high-performance mode, and the like. The data collector may modify the sensor collection routine based on a sensed change in a mode of operation, such as where the sensed change is a failure condition, a performance condition, a power condition, a temperature condition, a vibration condition, and the like. The evaluation of the received output data may be based on a collection routine with respect to a collection parameter, wherein the parameter is a network availability, a sensor availability, a time-based collection routine (e.g., collects sensor data on a schedule or over time), and the like.

[0898] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a collector route template, sensor specifications for sensors that correspond to the input channels, wherein the collector route template comprises a sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and providing a data analysis circuit structured to receive output data from the plurality of input channels and evaluate the received output data with respect to a rule, wherein the data collector is configured to modify the sensor collection routine based on the application of the rule to the received output data. In example embodiments, the method may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment.

[0899] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a collector route template, sensor specifications for sensors that correspond to the input channels, wherein the collector route template comprises a sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels; and providing a data analysis circuit structured to receive output data from the plurality of input channels and evaluate the received output data with respect to a

rule, wherein the data collector is configured to modify the sensor collection routine based on the application of the rule to the received output data. In example embodiments, the instructions may be deployed locally on the data collector, such as deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, where each of the input channels correspond to a sensor located in the environment.

[0900] Rapid route creation and modification in an industrial environment may employ smart route changes based on incoming data or alarms, such as changes enabling dynamic selection of data collection for analysis or correlation. Smart route changes may enable the system to alter current routing of sensor data based on incoming data or alarms. For instance, a user may set up a routing configuration that establishes a schedule of sensor collection for analysis, but when the analysis (or an alarm) indicates a special need, the system may change the sensor routing to address that need. For example, in the case where a change in a motor vibration profile (as one example among any of the machines described throughout this disclosure), such as rapidly increasing the peak amplitude of shaking on at least one axis of a vibration sensor set, that indicates a potential early failure of the motor, the system may change the routing to collect more focused data collection for analysis, such as initiating collection on more axes of the motor, initiating collection on additional bearings of the motor, and/or initiating collection using other sensors (such as temperature or heat flux sensors), that may confirm an initial hypothesis that the failure mode is occurring or otherwise assist in analysis of the state or operational condition of the machine.

[0901] Detected operational mode changes may trigger a rapid route change. For instance, an operational mode may be detected as the result of a single-point sensor out-of-range detection, an analysis determination, and the like, and generate a routing change. An analysis determination may be detected from a sensor end-point, such as through a single-point sensor analysis, a multiple-point sensor analysis, an analysis domain analysis (e.g., through a time profile, frequency profile, correlated multi-point determination), and the like. In another instance, a maintenance mode may be detected during routine maintenance, where a routing change increases data collection to capture data at a higher rate under an anomalous condition. A failure mode may be detected, such as through an alarm that indicates near-term potential for a failure of a machine that triggers increased data capture rate for analysis. Performance-based modes may be detected, such as detecting a level of output rate (e.g., peak, slack, idle), which may then initiate changes in routing to accommodate the analysis needs for the different performance monitoring and metrics associated with the state. For example, if a high peak speed is detected for a motor, a conveyor, an assembly line, a generator, a turbine, or the like, relative to historical measurements over some time period, additional sensors may be engaged to watch for failures that are typically associated with peak speeds, such as overheating (as measured by engaging a temperature or heat flux sensor), excessive noise (as measured by an acoustic or noise sensor), excessive shaking (as measured by one or more vibration sensors), or the like.

[0902] Alarm detections may trigger a rapid route change. Alarm sources may include a front-end collector, local intelligence resource, back-end data analysis process, ambient environment detector, network quality detector, power quality detector, heat, smoke, noise, flooding, and the

like. Alarm types may include a single-instance anomaly detection, multiple-instance anomaly detection, simultaneous multi-sensor detection, time-clustered sensor detection (e.g., a single sensor or multiple sensors), frequency-profile detection (e.g., increasing rate of anomaly detection such as an alarm increasing in its occurrence over time, a change in a frequency component of a sensor output such as a motor's physical vibration profile changing over time), and the like.

[0903] A machine learning system may change routing based on learned alarm pattern analysis. The machine learning system may learn system alarm condition patterns, such as alarm conditions expected under normal operating conditions, under peak operating conditions, expected over time based on age of components (e.g., new, during operational life, during extended life, during a warrantee period), and the like. The machine learning system may change routing based on a change in an alarm pattern, such as a system operating normally but experiencing a peak operating alarm pattern (e.g., a system running when it should not be), a system is new but experiencing an older profile (e.g., detection of infant mortality), and the like. The machine learning system may change routing based on a current alarm profile vs. an expected change in production condition. For example, a plant, system, or component is experiencing above average alarm conditions just before a ramp-up of production (e.g., could be foretelling of above average failures during increased production), just before going slack (e.g., could be an opportunity to ramp up maintenance procedures based on increased data taking routing scheme), after an unplanned event (e.g., weather, power outage, restart), and the like.

[0904] A rapid route change action may include: an increased rate of sampling (e.g., to a single sensor, to multiple sensors), an increase in the number of sensors being sampled (e.g., simultaneous sampling of other sensors on a device, coordinated sampling of similar sensors on near-by devices), generating a burst of sampling (e.g., sampling at a high rate for a period of time), and the like. Actions may be executed on a schedule, coordinated with a trigger, based on an operational mode, and the like. Triggered actions may include: anomalous data, an exceeded threshold level, an operational event trigger (e.g., at startup condition such as for startup motor torque), and the like.

[0905] A rapid route change may switch between routing schemes, such as an operational routing scheme (e.g., a subset of sensor collection for normal operations), a scheduled maintenance routing scheme (e.g., an increased and focused set of sensor collection than for normal operations), and the like. The distribution of sensor data may be changed, such as to distribute sensor collection across the system, such as for a sensor collection set for specific components, functions, and modes. A failure mode routing scheme may entail multiple focused sensor collection groups targeting different failure mode analyses (e.g., for a motor, one failure mode may be for bearings, another for startup speed-torque) where a different subset of sensor data may be needed based on the failure mode (e.g., as detected in anomalous readings taken during operations or maintenance). Power saving mode routing may be executed when weather conditions necessitate reduced plant power.

[0906] Dynamic adjustment of route changes may be executed based on connectivity factors, such as the factors associated with the collector or network availability and bandwidth. For example, routing may be changed for a device associated with an alarm detection, where changing routing for targeted devices on the network frees up bandwidth. Changes to routing may have a duration,

such as only for a pre-determined period of time and then switching back, maintaining a change until user-directed, changing duration based on network availability, and the like.

[0907] In example embodiments, and referring to Figure 101, smart route changes may be implemented by a local data collection system 10520 for collection and monitoring of data collected through a plurality of input channels 10500, such as data from sensors 10522, IoT devices 10524, and the like. The local collection system 102, also referred to herein as a data collector 10520, may comprise a data storage 10502, a data acquisition circuit 10504, a data analysis circuit 10506, a response circuit 10508, and the like, wherein the monitoring facilities may be deployed locally on the data collector 10520, in part locally on the data collector and in part on a remote information technology infrastructure component apart from the data collector, and the like. Smart route changes may be implemented between data collectors, such as where a state message is transmitted between the data collectors (e.g., from an input channel that is mounted in proximity to a second input channel, from a related group of input sensors, and the like). A monitoring system may comprise a plurality of input channels 10500 communicatively coupled to the data collector 10520. The data acquisition circuit 10504 may be structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels 10500, wherein the data acquisition circuit 10504 acquires sensor data from a first route of input channels for the plurality of input channels. The data storage 10502 may be structured to store sensor data, sensor specifications, and the like, for sensors 10524 that correspond to the input channels 10500. The data analysis circuit 10506 may be structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information may include an alarm threshold level, and wherein the data analysis circuit 10506 sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels. Further, the data analysis circuit 10506 may transmit the alarm state across a network to a routing control facility 10512. The response circuit 10508 may be structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels upon reception of a routing change indication from the routing control facility. In the case of a network transmission, the alternate routing of input channels may include the first input channel and a group of input channels related to the first input channel, where the data collector executes the change in routing of the input channels if a communication parameter of the network between the data collector and the routing control facility is not met (e.g., a time-period parameter, a network connection and/or bandwidth availability parameter).

[0908] In example embodiments, an alarm state may indicate a detection mode, such as an operational mode detection comprising an out-of-range detection, a maintenance mode detection comprising an alarm detected during maintenance, a failure mode detection (e.g., where the controller communicates a failure mode detection facility), a power mode detection wherein the alarm state is indicative of a power related limitation data of the anticipated state information, a performance mode detection wherein the alarm state is indicative of a high-performance limitation data of the anticipated state information, and the like. The monitoring system may further include the analysis circuit setting the alarm state when the alarm threshold level is exceeded for an

alternate input channel in the first group of input channels, such as where the setting of the alarm state for the first input channel and the alternate input channel are determined to be a multiple-instance anomaly detection, wherein the second routing of input channels comprises the first input channel and a second input channel, wherein the sensor data from the first input channel and the second input channel contribute to simultaneous data analysis. The second routing of input channels may include a change in a routing collection parameter, such as where the routing collection parameter is an increase in sampling rate, an increase in the number of channels being sampled, a burst sampling of at least one of the plurality of input channels, and the like.

[0909] In example embodiments, and referring to Figure 100, collector route templates 10510 may be utilized for smart route changes and may be implemented by a local data collection system 10512 for collection and monitoring of data collected through a plurality of input channels 10500, such as data from sensors 10514, IoT devices 10516, and the like. The local collection system 10512, also referred to herein as a data collector 10512, may comprise a data storage 10502, a data acquisition circuit 10504, a data analysis circuit 10506, a response circuit 10508, and the like, wherein the monitoring facilities may be deployed locally on the data collector 10512, in part locally on the data collector and in part on a remote information technology infrastructure component apart from the data collector, and the like.

[0910] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collector communicatively coupled to a plurality of input channels; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; a data storage structured to store sensor specifications for sensors that correspond to the input channels; a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels; and a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels, wherein the alternate routing of input channels comprise the first input channel and a group of input channels related to the first input channel. In example embodiments, the system may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment. The group of input channels may be related to the first input channel are at least in part taken from the plurality of input channels not included in the first routing of input channels. An alarm state may indicate a detection mode, such as where the detection mode is an operational mode detection comprising an out-of-range detection, the detection mode is a maintenance mode detection comprising an alarm detected during maintenance, the detection mode is a failure mode detection. The controller may communicate the failure mode detection facility, such as where the detection mode is a power mode detection and the alarm state is indicative of a

power related limitation data of the anticipated state information, the detection mode is a performance mode detection and the alarm state is indicative of a high-performance limitation data of the anticipated state information, and the like. The analysis circuit may set the alarm state when the alarm threshold level is exceeded for an alternate input channel in the first group of input channels, such as where the setting of the alarm state for the first input channel and the alternate input channel are determined to be a multiple-instance anomaly detection, wherein the alternate routing of input channels comprises the first input channel and a second input channel, wherein the sensor data from the first input channel and the second input channel contribute to simultaneous data analysis. The alternate routing of input channels may include a change in a routing collection parameter, such as for an increase in sampling rate, an increase in the number of channels being sampled, a burst sampling of at least one of the plurality of input channels, and the like.

[0911] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels; and providing a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels, wherein the alternate routing of input channels comprise the first input channel and a group of input channels related to the first input channel. In example embodiments, the system may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0912] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels;

and providing a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels, wherein the alternate routing of input channels comprise the first input channel and a group of input channels related to the first input channel. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0913] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collector communicatively coupled to a plurality of input channels; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; a data storage structured to store sensor specifications for sensors that correspond to the input channels; a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels and transmits the alarm state across a network to a routing control facility; and a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels upon reception of a routing change indication from the routing control facility, wherein the alternate routing of input channels comprise the first input channel and a group of input channels related to the first input channel, wherein the data collector automatically executes the change in routing of the input channels if a communication parameter of the network between the data collector and the routing control facility is not met. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment. The communication parameter may be a time-period parameter within which the routing control facility must respond. The communication parameter may be a network availability parameter, such as a network connection parameter or bandwidth requirement. The group of input channels related to the first input channel may be at least in part taken from the plurality of input channels not included in the first routing of input channels. The alarm state may indicate a detection mode, such as an operational mode detection comprising an out-of-range detection, a maintenance mode detection comprising an alarm detected during maintenance, and the like. The detection mode may be a failure mode detection, such as when the controller communicates the failure mode detection facility, the alarm state is indicative of a power related limitation data of the anticipated state information, the detection mode is a performance mode detection where the alarm state is indicative of a high-performance limitation data of the anticipated state information, and the like. The analysis circuit may set the alarm state when the alarm threshold level is exceeded for an alternate input channel in the first group of input channels, such as where the setting of the alarm

state for the first input channel and the alternate input channel are determined to be a multiple-instance anomaly detection, wherein the alternate routing of input channels comprises the first input channel and a second input channel, wherein the sensor data from the first input channel and the second input channel contribute to simultaneous data analysis. The alternate routing of input channels may be a change in a routing collection parameter, such as an increase in sampling rate, is an increase in the number of channels being sampled, a burst sampling of at least one of the plurality of input channels, and the like.

[0914] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels and transmits the alarm state across a network to a routing control facility; and providing a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels upon reception of a routing change indication from the routing control facility, wherein the alternate routing of input channels comprise the first input channel and a group of input channels related to the first input channel, wherein the data collector automatically executes the change in routing of the input channels if a communication parameter of the network between the data collector and the routing control facility is not met. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0915] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels

and transmits the alarm state across a network to a routing control facility; and providing a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels upon reception of a routing change indication from the routing control facility, wherein the alternate routing of input channels comprise the first input channel and a group of input channels related to the first input channel, wherein the data collector automatically executes the change in routing of the input channels if a communication parameter of the network between the data collector and the routing control facility is not met. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0916] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a first and second data collector communicatively coupled to a plurality of input channels; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; a data storage structured to store sensor specifications for sensors that correspond to the input channels; a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels; a communication circuit structured to communicate with a second data collector, wherein the second data collector transmits a state message related to a first input channel from the first route of input channels; and a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels based on the state message from the second data collector, wherein the alternate routing of input channel comprise the first input channel and a group of input channels related to the first input sensor. In example embodiments, the system may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment. The set state message transmitted from the second data collector may be from a second input channel that is mounted in proximity to the first input channel. The set alarm transmitted from the second controller may be from a second input sensor that is part of a related group of input sensors comprising the first input sensor. The group of input channels related to the first input channel may be at least in part taken from the plurality of input channels not included in the first routing of input channels. The alarm state may indicate a detection mode, such as where the detection mode is an operational mode detection comprising an out-of-range detection, a maintenance mode detection comprising an alarm detected during maintenance, is a failure mode detection, and the like. The controller may communicate the failure mode detection facility, such as where the detection mode is a power mode detection and the alarm state is

indicative of a power related limitation data of the anticipated state information, the detection mode is a performance mode detection where the alarm state is indicative of a high-performance limitation data of the anticipated state information, and the like. The analysis circuit may set the alarm state when the alarm threshold level is exceeded for an alternate input channel in the first group of input channels, such as where the setting of the alarm state for the first input channel and the alternate input channel are determined to be a multiple-instance anomaly detection, wherein the alternate routing of input channels comprises the first input channel and a second input channel, wherein the sensor data from the first input channel and the second input channel contribute to simultaneous data analysis. The alternate routing of input channels may be a change in a routing collection parameter, such as an increase in sampling rate, an increase in the number of channels being sampled, a burst sampling of at least one of the plurality of input channels, and the like.

[0917] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a first and second data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels; providing a communication circuit structured to communicate with a second data collector, wherein the second data collector transmits a state message related to a first input channel from the first route of input channels, and providing a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels based on the state message from the second data collector, wherein the alternate routing of input channel comprise the first input channel and a group of input channels related to the first input sensor. In example embodiments, the method may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0918] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a first and second data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels for the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit

structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels; providing a communication circuit structured to communicate with a second data collector, wherein the second data collector transmits a state message related to a first input channel from the first route of input channels, and providing a response circuit structured to change the routing of the input channels for data collection from the first routing of input channels to an alternate routing of input channels based on the state message from the second data collector, wherein the alternate routing of input channel comprise the first input channel and a group of input channels related to the first input sensor. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0919] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collector communicatively coupled to a plurality of input channels; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channel, wherein the data acquisition circuit acquires sensor data from a first group of input channels from the plurality of input channels; a data storage structured to store sensor specifications for sensors that correspond to the input channels; a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channel; and a response circuit structured to change the input channels being collected from the first group of input channels to an alternative group of input channels, wherein the alternate group of input channels comprise the first input channel and a group of input channels related to the first input sensor. In example embodiments, the system may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment. The group of input sensors related to the first input sensor may be at least in part taken from the plurality of input sensors not included in the first group of input sensors. The first group of input channels related to the first input channel may be at least in part taken from the plurality of input channels not included in the first routing of input channels. The alarm state may indicate a detection mode, such as where the detection mode is an operational mode detection comprising an out-of-range detection, a maintenance mode detection comprising an alarm detected during maintenance. The detection mode may be a failure mode detection, such as where the controller communicates the failure mode detection facility. The detection mode may be a power mode detection where the alarm state is indicative of a power related limitation data of the anticipated state information. The detection mode may be a performance mode detection, where the alarm state is indicative of a high-performance limitation data of the anticipated state

information. The analysis circuit may set the alarm state when the alarm threshold level is exceeded for an alternate input channel in the first group of input channels, such as when the setting of the alarm state for the first input channel and the alternate input channel are determined to be a multiple-instance anomaly detection, wherein the alternate routing of input channels comprises the first input channel and a second input channel, wherein the sensor data from the first input channel and the second input channel contribute to simultaneous data analysis. An alternative group of input channels may include a change in a routing collection parameter, such as where the routing collection parameter is an increase in sampling rate, an increase in the number of channels being sampled, a burst sampling of at least one of the plurality of input channels, and the like.

[0920] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channel, wherein the data acquisition circuit acquires sensor data from a first group of input channels from the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channel; and providing a response circuit structured to change the input channels being collected from the first group of input channels to an alternative group of input channels, wherein the alternate group of input channels comprise the first input channel and a group of input channels related to the first input sensor. In example embodiments, the method may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0921] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a data collector communicatively coupled to a plurality of input channels; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channel, wherein the data acquisition circuit acquires sensor data from a first group of input channels from the plurality of input channels; providing a data storage structured to store sensor specifications for sensors that correspond to the input channels; providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channel; and providing a response circuit structured to change the input channels being collected from the first group of input channels to an alternative group of input channels, wherein

the alternate group of input channels comprise the first input channel and a group of input channels related to the first input sensor. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0922] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collector communicatively coupled to a plurality of input channels; a data storage structured to store a plurality of collector route templates, sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels; and a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels, wherein the data collector is configured to switch from a current routing template collection routine to an alternate routing template collection routine based on a setting of an alarm state. In example embodiments, the system may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment. The setting of the alarm state may be based on operational mode routing collection schemes, such as where the operational mode is at least one of a normal operational mode, a peak operational mode, an idle operational mode, a maintenance operational mode, and a power saving operational mode. The alarm threshold level may be associated with a sensed change to one of the plurality of input channels, such as where the sensed change is a failure condition, is a performance condition, a power condition, a temperature condition, a vibration condition, and the like. The alarm state may indicate a detection mode, such as where the detection mode is an operational mode detection comprising an out-of-range detection, a maintenance mode detection comprising an alarm detected during maintenance, and the like. The detection mode may be a power mode detection, where the alarm state is indicative of a power related limitation data of the anticipated state information. The detection mode may be a performance mode detection, where the alarm state is indicative of a high-performance limitation data of the anticipated state information. The analysis circuit may set the alarm state when the alarm threshold level is exceeded for an alternate input channel, such as wherein the setting of the alarm state is determined to be a multiple-instance anomaly detection. The alternate routing template may be a change to an input channel routing collection parameter. The routing collection parameter may be an increase in sampling rate, such as an increase in the number of channels being sampled, a burst sampling of at least one of the plurality of input channels, and the like.

[0923] In example embodiments, a computer-implemented method for implementing a monitoring system for data collection in an industrial environment may comprise: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a plurality of collector route templates, sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels; and providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels, wherein the data collector is configured to switch from a current routing template collection routine to an alternate routing template collection routine based on a setting of an alarm state. In example embodiments, the system may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0924] In example embodiments, one or more non-transitory computer-readable media comprising computer executable instructions that, when executed, may cause at least one processor to perform actions comprising: providing a data collector communicatively coupled to a plurality of input channels; providing a data storage structured to store a plurality of collector route templates, sensor specifications for sensors that correspond to the input channels, wherein the plurality of collector route templates each comprise a different sensor collection routine; providing a data acquisition circuit structured to interpret a plurality of detection values, each of the plurality of detection values corresponding to at least one of the input channels, wherein the data acquisition circuit acquires sensor data from a first route of input channels; and providing a data analysis circuit structured to evaluate the sensor data with respect to stored anticipated state information, wherein the anticipated state information comprises an alarm threshold level, and wherein the data analysis circuit sets an alarm state when the alarm threshold level is exceeded for a first input channel in the first group of input channels, wherein the data collector is configured to switch from a current routing template collection routine to an alternate routing template collection routine based on a setting of an alarm state. In example embodiments, the instructions may be deployed locally on the data collector, deployed in part locally on the data collector and in part on a remote information technology infrastructure component apart from the collector, wherein each of the input channels correspond to a sensor located in the environment.

[0925] Methods and systems are disclosed herein for a system for data collection in an industrial environment using intelligent management of data collection bands, referred to herein in some cases as smart bands. Smart bands may facilitate intelligent, situational, context-aware collection of data, such as by a data collector (such as any of the wide range of data collector embodiments

described throughout this disclosure). Intelligent management of data collection via smart bands may improve various parameters of data collection, as well as parameters of the processes, applications, and products that depend on data collection, such as data quality parameters, consistency parameters, efficiency parameters, comprehensiveness parameters, reliability parameters, effectiveness parameters, storage utilization parameters, yield parameters (including financial yield, output yield, and reduction of adverse events), energy consumption parameters, bandwidth utilization parameters, input/output speed parameters, redundancy parameters, security parameters, safety parameters, interference parameters, signal-to-noise parameters, statistical relevancy parameters, and others. Intelligent management of smart bands may optimize across one or more such parameters, such as based on a weighting of the value of the parameters; for example, a smart band may be managed to provide a given level of redundancy for critical data, while not exceeding a specified level of energy usage. This may include using a variety of optimization techniques described throughout this disclosure and the documents incorporated herein by reference.

[0926] In example embodiments, such methods and systems for intelligent management of smart bands include an expert system and supporting technology components, services, processes, modules, applications and interfaces, for managing the smart bands (collectively referred to in some cases as a smart band platform 10722), which may include a model-based expert system, a rule-based expert system, an expert system using artificial intelligence (such as a machine learning system, which may include a neural net expert system, a self-organizing map system, a human-supervised machine learning system, a state determination system, a classification system, or other artificial intelligence system), or various hybrids or combinations of any of the above. References to an expert system should be understood to encompass utilization of any one of the foregoing or suitable combinations, except where context indicates otherwise. Intelligent management may be of data collection of various types of data (e.g., vibration data, noise data and other sensor data of the types described throughout this disclosure) for event detection, state detection, and the like. Intelligent management may include managing a plurality of smart bands each directed at supporting an identified application, process or workflow, such as confirming progress toward or alignment with one or more objectives, goals, rules, policies, or guidelines. Intelligent management may also involve managing data collection bands targeted to backing out an unknown variable based on collection of other data (such as based on a model of the behavior of a system that involves the variable), selecting preferred inputs among available inputs (including specifying combinations, fusions, or multiplexing of inputs), and/or specifying an input band among available input bands.

[0927] Data collection bands, or smart bands, may include any number of items such as sensors, input channels, data locations, data streams, data protocols, data extraction techniques, data transformation techniques, data loading techniques, data types, frequency of sampling, placement of sensors, static data points, metadata, fusion of data, multiplexing of data, and the like as described herein. Smart band settings, which may be used interchangeably with smart band and data collection band, may describe the configuration and makeup of the smart band, such as by

specifying the parameters that define the smart band. For example, data collection bands, or smart bands, may include one or more frequencies to measure. Frequency data may further include at least one of a group of spectral peaks, a true-peak level, a crest factor derived from a time waveform, and an overall waveform derived from a vibration envelope, as well as other signal characteristics described throughout this disclosure. Smart bands may include sensors measuring or data regarding one or more wavelengths, one or more spectra, and/or one or more types of data from various sensors and metadata. Smart bands may include one or more sensors or types of sensors of a wide range of types, such as described throughout this disclosure and the documents incorporated by reference herein. Indeed, the sensors described herein may be used in any of the methods or systems described throughout this disclosure. For example, one sensor may be an accelerometer, such as one that measures voltage per G ("V/G") of acceleration (e.g., 100 mV/G, 500 mV/G, 1 V/G, 5 V/G, 10 V/G, and the like). In example embodiments, the data collection band circuit may alter the makeup of the subset of the plurality of sensors used in a smart band based on optimizing the responsiveness of the sensor, such as for example choosing an accelerometer better suited for measuring acceleration of a low speed mixer versus one better suited for measuring acceleration of a high speed industrial centrifuge. Choosing may be done intelligently, such as for example with a proximity probe and multiple accelerometers disposed on a centrifuge where while at low speed, one accelerometer is used for measuring in the smart band and another is used at high speeds. Accelerometers come in various types, such as piezo-electric crystal, low frequency (e.g., 10 V/G), high speed compressors (10 MV/G), MEMS, and the like. In another example, one sensor may be a proximity probe which can be used for sleeve or tilt-pad bearings (e.g., oil bath), or a velocity probe. In yet another example, one sensor may be a solid-state relay (SSR) that is structured to automatically interface with a routed data collector (such as a mobile or portable data collector) to obtain or deliver data. In another example, a mobile or portable data collector may be routed to alter the makeup of the plurality of available sensors, such as by bringing an appropriate accelerometer to a point of sensing, such as on or near a component of a machine. In still another example, one sensor may be a triax probe (e.g., a 100 MV/G triax probe), that in embodiments is used for portable data collection. In example embodiments, of a triax probe, a vertical element on one axis of the probe may have a high frequency response while the ones mounted horizontally may influence the frequency response of the whole triax. In another example, one sensor may be a temperature sensor and may include a probe with a temperature sensor built inside, such as to obtain a bearing temperature. In still additional examples, sensors may be ultrasonic, microphone, touch, capacitive, vibration, acoustic, pressure, strain gauges, thermographic (e.g., camera), imaging (e.g., camera, laser, IR, structured light), a field detector, an EMF meter to measure an AC electromagnetic field, a gaussmeter, a motion detector, a chemical detector, a gas detector, a CBRNE detector, a vibration transducer, a magnetometer, positional, location-based, a velocity sensor, a displacement sensor, a tachometer, a flow sensor, a level sensor, a proximity sensor, a pH sensor, a hygrometer / moisture sensor, a densitometric sensor, an anemometer, a viscometer, or any analog industrial sensor and/or digital industrial sensor. In a further example, sensors may be directed at detecting or measuring ambient noise, such as a sound sensor or microphone, an

ultrasound sensor, an acoustic wave sensor, and an optical vibration sensor (e.g., using a camera to see oscillations that produce noise). In still another example, one sensor may be a motion detector.

[0928] Data collection bands, or smart bands, may be of or may be configured to encompass one or more frequencies, wavelengths, or spectra for particular sensors, for particular groups of sensors, or for combined signals from multiple sensors (such as involving multiplexing or sensor fusion).

[0929] Data collection bands, or smart bands, may be of or may be configured to encompass one or more sensors or sensor data (including groups of sensors and combined signals) from one or more pieces of equipment/components, areas of an installation, disparate but interconnected areas of an installation (e.g., a machine assembly line and a boiler room used to power the line), or locations (e.g., a building in Cambridge and a building in Boston). Smart band settings, configurations, instructions, or specifications (collectively referred to herein using any one of those terms) may include where to place a sensor, how frequently to sample a data point or points, the granularity at which a sample is taken (e.g., a number of sampling points per fraction of a second), which sensor of a set of redundant sensors to sample, an average sampling protocol for redundant sensors, and any other aspect that would affect data acquisition.

[0930] Within the smart band platform 10722, an expert system, which may comprise a neural net, a model-based system, a rule-based system, a machine learning data analysis circuit, and/or a hybrid of any of those, may begin iteration towards convergence on a smart band that is optimized for a particular goal or outcome, such as predicting and managing performance, health, or other characteristics of a piece of equipment, a component, or a system of equipment or components. Based on continuous or periodic analysis of sensor data, as patterns/trends are identified, or outliers appear, or a group of sensor readings begin to change, etc., the expert system may modify its data collection bands intelligently. This may occur by triggering a rule that reflects a model or understanding of system behavior (e.g., recognizing a shift in operating mode that calls for different sensors as velocity of a shaft increases) or it may occur under control of a neural net (either in combination with a rule-based approach or on its own), where inputs are provided such that the neural net over time learns to select appropriate collection modes based on feedback as to successful outcomes (e.g., successful classification of the state of a system, successful prediction, successful operation relative to a metric, or the like). For example, when a new pressure reactor is installed in a chemical processing facility, data from the current data collection band may not accurately predict the state or metric of operation of the system, thus, the machine learning data analysis circuit may begin to iterate to determine if a new data collection band is better at predicting a state. Based on offset system data, such as from a library or other data structure, certain sensors, frequency bands or other smart band members may be used in the smart band initially and data may be collected to assess performance. As the neural net iterates, other sensors/frequency bands may be accessed to determine their relative weight in identifying performance metrics. Over time, a new frequency band may be identified (or a new collection of sensors, a new set of configurations for sensors, or the like) as a better gauge of performance in the system and the expert system may modify its data collection band based on this iteration. For example, perhaps a slightly different or

older associated turbine agitator in a chemical reaction facility dampens one or more vibration frequencies while a different frequency is of higher amplitude and present during optimal performance than what was seen in the offset system. In this example, the smart band may be altered from what was suggested by the corresponding offset system to capture the higher amplitude frequency that is present in the current system.

[0931] The expert system, in embodiments involving a neural net or other machine learning system, may be seeded and may iterate, such as towards convergence on a smart band, based on feedback and operation parameters, such as described herein. Certain feedback may include utilization measures, efficiency measures (e.g., power or energy utilization, use of storage, use of bandwidth, use of input/output use of perishable materials, use of fuel, and/or financial efficiency), measures of success in prediction or anticipation of states (e.g., avoidance and mitigation of faults), productivity measures (e.g., workflow), yield measures, and profit measures. Certain parameters may include: storage parameters (e.g., data storage, fuel storage, storage of inventory and the like); network parameters (e.g., network bandwidth, input/output speeds, network utilization, network cost, network speed, network availability and the like); transmission parameters (e.g., quality of transmission of data, speed of transmission of data, error rates in transmission, cost of transmission and the like); security parameters (e.g., number and/or type of exposure events; vulnerability to attack, data loss, data breach, access parameters, and the like); location and positioning parameters (e.g., location of data collectors, location of workers, location of machines and equipment, location of inventory units, location of parts and materials, location of network access points, location of ingress and egress points, location of landing positions, location of sensor sets, location of network infrastructure, location of power sources and the like); input selection parameters, data combination parameters (e.g., for multiplexing, extraction, transformation, loading, and the like); power parameters; states (e.g., operating modes, availability states, environmental states, fault modes, maintenance modes, anticipated states); events; and equipment specifications. With respect to states, operating modes may include mobility modes (direction, speed, acceleration, and the like), type of mobility modes (e.g., rolling, flying, sliding, levitation, hovering, floating, and the like), performance modes (e.g., gears, rotational speeds, heat levels, assembly line speeds, voltage levels, frequency levels, and the like), output modes, fuel conversion modes, resource consumption modes, and financial performance modes (e.g., yield, profitability, and the like). Availability states may refer to anticipating conditions that could cause a machine to go offline or require backup. Environmental states may refer to ambient temperature, ambient humidity/moisture, ambient pressure, ambient wind/fluid flow, presence of pollution or contaminants, presence of interfering elements (e.g., electrical noise, vibration), power availability, and power quality. Anticipated states may include: achieving or not achieving a desired goal, such as a specified/threshold output production rate, a specified/threshold generation rate, an operational efficiency/ failure rate, a financial efficiency/profit goal, a power efficiency/resource utilization; an avoidance of a fault condition (e.g., overheating, slow performance, excessive speed, excessive motion, excessive vibration/oscillation, excessive acceleration, expansion/contraction, electrical failure, running out of stored power/fuel, overpressure, excessive radiation/melt down, fire, freezing, failure of fluid

flow (e.g., stuck valves, frozen fluids); mechanical failures (e.g., broken component, worn component, faulty coupling, misalignment, asymmetries/deflection, damaged component (e.g., deflection, strain, stress, cracking], imbalances, collisions, jammed elements, and lost or slipping chain or belt); avoidance of a dangerous condition or catastrophic failure; and availability (online status).

[0932] The expert system may comprise or be seeded with a model that predicts an outcome or state given a set of data (which may comprise inputs from sensors, such as via a data collector, as well as other data, such as from system components, from external systems and from external data sources). For example, the model may be an operating model for an industrial environment, machine, or workflow. In another example, the model may be for anticipating states, for predicting fault and optimizing maintenance, for self-organizing storage (e.g., on devices, in data pools and/or in the cloud), for optimizing data transport (such as for optimizing network coding, network-condition-sensitive routing, and the like), for optimizing data marketplaces, and the like.

[0933] The iteration of the expert system may result in any number of downstream actions based on analysis of data from the smart band. In example embodiments, the expert system may determine that the system should either keep or modify operational parameters, equipment or a weighting of a neural net model given a desired goal, such as a specified/threshold output production rate, specified/threshold generation rate, an operational efficiency/failure rate, a financial efficiency/profit goal, a power efficiency/resource utilization, an avoidance of a fault condition, an avoidance of a dangerous condition or catastrophic failure, and the like. In example embodiments, the adjustments may be based on determining context of an industrial system, such as understanding a type of equipment, its purpose, its typical operating modes, the functional specifications for the equipment, the relationship of the equipment to other features of the environment (including any other systems that provide input to or take input from the equipment), the presence and role of operators (including humans and automated control systems), and ambient or environmental conditions. For example, in order to achieve a profit goal, a pipeline in a refinery may need to operate for a certain amount of time a day and/or at a certain flow rate. The expert system may be seeded with a model for operation of the pipeline in a manner that results in a specified profit goal, such as indicating a given flow rate of material through the pipeline based on the current market sale price for the material and the cost of getting the material into the pipeline. As it acquires data and iterates, the model will predict whether the profit goal will be achieved given the current data. Based on the results of the iteration of the expert system, a recommendation may be made (or a control instruction may be automatically provided) to operate the pipeline at a higher flow rate, to keep it operational for longer or the like. Further, as the system iterates, one or more additional sensors may be sampled in the model to determine if their addition to the smart band would improve predicting a state. In another embodiment, the expert system may determine that the system should either keep or modify operational parameters, equipment or a weighting of a neural net or other model given a constraint of operation (e.g., meeting a required endpoint (e.g., delivery date, amount, cost, coordination with another system), operating with a limited resource (e.g., power, fuel, battery), storage (e.g., data storage), bandwidth (e.g., local network, p2p, WAN,

internet bandwidth, availability, or input/output capacity), authorization (e.g., role-based)), a warranty limitation, a manufacturer's guideline, a maintenance guideline). For example, a constraint of operating a boiler in a refinery is that the aeration of the boiler feedwater needs to be reduced in the cycle; therefore, the boiler must coordinate with the deaerator. In this example, the expert system is seeded with a model for operation of the boiler in coordination with the de-aerator that results in a specified overall performance. As sensor data from the system is acquired, the expert system may determine that an aspect of one or both of the boiler and aerator must be changed to continue to achieve the specific overall performance. In a further embodiment, the expert system may determine that the system should either keep or modify operational parameters, equipment or a weighting of a neural net model given an identified choke point. In still another embodiment, the expert system may determine that the system should either keep or modify operational parameters, equipment or a weighting of a neural net model given an off-nominal operation. For example, a reciprocating compressor in a refinery that delivers gases at high pressure may be measured as having an off-nominal operation by sensors that feed their data into an expert system (optionally including a neural net or other machine learning system). As the expert system iterates and receives the off-nominal data, it may predict that the refinery will not achieve a specified goal and will recommend an action, such as taking the reciprocating compressor offline for maintenance. In another embodiment, the expert system may determine that the system should collect more/fewer data points from one or more sensors. For example, an anchor agitator in a pharmaceutical processing plant may be programmed to agitate the contents of a tank until a certain level of viscosity (e.g., as measured in centipoise) is obtained. As the expert system collects data throughout the run indicating an increase in viscosity, the expert system may recommend collecting additional data points to confirm a predicted state in the face of the increased strain on the plant systems from the viscosity. In yet another embodiment, the expert system may determine that the system should change a data storage technique. In still another example, the expert system may determine that the system should change a data presentation mode or manner. In a further embodiment, the expert system may determine that the system should apply one or more filters (low pass, high pass, band pass, etc.) to collected data. In yet a further embodiment, the expert system may determine that the system should collect data from a new smart band/new set of sensors and/or begin measuring a new aspect that the neural net identified itself. For example, various measurements may be made of paddle-type agitator mixers operating in a pharmaceutical plant, such as mixing times, temperature, homogeneous substrate distribution, heat exchange with internal structures and the tank wall or oxygen transfer rate, mechanical stress, forces and torques on agitator vessels and internal structures, and the like. Various sensor data streams may be included in a smart band monitoring these various aspects of the paddle-type agitator mixer, such as a flow meter, a thermometer, and others. As the expert system iterates, perhaps having been seeded with minimal data from during the agitator's run, a new aspect of the operation may become apparent, such as the impact of pH on the state of the run. Thus, a new smart band will be identified by the expert system that includes sensor data from a pH meter. In yet still a further embodiment, the expert system may determine that the system should discontinue collection of data from a smart

band, one or more sensors, or the like. In another embodiment, the expert system may determine that the system should initiate data collection from a new smart band, such as a new smart band identified by the neural net itself. In yet another embodiment, the expert system may determine that the system should adjust the weights/biases of a model used by the expert system. In still another embodiment, the expert system may determine that the system should remove/re-task under-utilized equipment. For example, a plurality of agitators working with a pump blasting liquid in a pharmaceutical processing plant may be monitored during operation of the plant by the expert system. Through iteration of the expert system seeded with data from a run of the plant with the agitators, the expert system may predict that a state will be achieved even if one or more agitators are taken out of service.

[0934] In example embodiments, a monitoring system for data collection in an industrial environment may include a plurality of input sensors, such as any of those described herein, communicatively coupled to a data collector having a controller. The monitoring system may include a data collection band circuit structured to determine at least one subset of the plurality of sensors from which to process output data. The monitoring system may also include a machine learning data analysis circuit structured to receive output data from the at least one subset of the plurality of sensors and learn received output data patterns indicative of a state. In example embodiments, the data collection band circuit may alter the at least one subset of the plurality of sensors, or an aspect thereof, based on one or more of the learned received output data patterns and the state. In certain example embodiments, the machine learning data analysis circuit is seeded with a model that enables it to learn data patterns. The model may be a physical model, an operational model, a system model, and the like. In other embodiments, the machine learning data analysis circuit is structured for deep learning wherein input data is fed to the circuit with no or minimal seeding and the machine learning data analysis circuit learns based on output feedback. For example, a static mixer in a chemical processing plant producing polymers may be used to facilitate the polymerization reaction. The static mixer may employ turbulent or laminar flow in its operation. Minimal data, such as heat transfer, velocity of flow out of the mixer, Reynolds number or pressure drop, acquired during the operation of the static mixer may be fed into the expert system which may iterate towards a prediction based on initial feedback (e.g., viscosity of the polymer, color of the polymer, reactivity of the polymer).

[0935] There may be a balance of multiple goals/guidelines in the management of smart bands by the expert system. For example, a repair and maintenance organization (RMO) may have operating parameters designed for maintenance of a storage tank in a refinery, while the owner of the refinery may have particular operating parameters for the storage tank that are designed for meeting a production goal. These goals, in this example relating to a maintenance goal or a production output, may be tracked by a different data collection bands. For example, maintenance of a storage tank may be tracked by sensors including a vibration transducer and a strain gauge, while the production goal of a storage tank may be tracked by sensors including a temperature sensor and a flow meter. The expert system may (optionally using a neural net, machine learning system, deep learning system, or the like, which may occur under supervision by one or more supervisors (human or

automated)) intelligently manage bands aligned with different goals and assign weights, parameter modifications, or recommendations based on a factor, such as a bias towards one goal or a compromise to allow better alignment with all goals being tracked, for example. Compromises among the goals delivered to the expert system may be based on one or more hierarchies or rules (relating to the authority, role, criticality, or the like) of the applicable goals. In example embodiments, compromises among goals may be optimized using machine learning, such as a neural net, deep learning system, or other artificial intelligence system as described throughout this disclosure. In one illustrative example, in a chemical processing plant where a gas-powered agitator is operating, the expert system may manage multiple smart bands, such as one directed to detecting the operational status of the gas-powered agitator, one directed at identifying a probability of hitting a production goal, and one directed at determining if the operation of the gas-powered agitator is meeting a fuel efficiency goal. Each of these smart bands may be populated with different sensors or data from different sensors (e.g., a vibration transducer to indicate operational status, a flow meter to indicate production goal, and a fuel gauge to indicate a fuel efficiency) whose output data are indicative of an aspect of the particular goal. Where a single sensor or a set of sensors is helpful for more than one goal, overlapping smart bands (having some sensors in common and other sensors not in common) may take input from that sensor or set of sensors, as managed by the smart band platform 10722. If there are constraints on data collection (such as due to power limitations, storage limitations, bandwidth limitations, input/output processing capabilities, or the like), a rule may indicate that one goal (e.g., a fuel utilization goal or a pollution reduction goal that is mandated by law or regulation) takes precedence, such that the data collection for the smart bands associated with that goal are maintained as others are paused or shut down. Management of prioritization of goals may be hierarchical or may occur by machine learning. The expert system may be seeded with models, or may not be seeded at all, in iterating towards a predicted state (i.e., meeting the goal) given the current data it has acquired. In this example, during operation of the gas-powered agitator, the plant owner may decide to bias the system towards fuel efficiency. All of the bands may still be monitored, but as the expert system iterates and predicts that the system will not meet or is not meeting a particular goal, and then offers recommended changes directed at increasing the chance of meeting the goal, the plant owner may structure the system with a bias towards fuel efficiency so that the recommended changes to parameters affecting fuel efficiency are made in favor of making other recommended changes.

[0936] In example embodiments, the expert system may continue iterating in a deep-learning fashion to arrive at a single smart band, after being seeded with more than one smart band, that optimizes meeting more than one goal. For example, there may be multiple goals tracked for a thermic heating system in a chemical processing or a food processing plant, such as thermal efficiency and economic efficiency. Thermal efficiency for the thermic heating system may be expressed by comparing BTUs put into the system, which can be obtained by knowing the amount of and quality of the fuel being used, and the BTUs out of the system, which is calculated using the flow out of the system and the temperature differential of materials in and out of the system. Economic efficiency of the thermic heating system may be expressed as the ratio between costs to

run the system (including fuel, labor, materials, and services) and energy output from the system for a period of time. Data used to track thermal efficiency may include data from a flow meter, quality data point(s), and a thermometer, and data used to track economic efficiency may be an energy output from the system (e.g., kWh) and costs data. These data may be used in smart bands by the expert system to predict states, however, the expert system may iterate toward a smart band that is optimized to predict states related to both thermal and economic efficiency. The new smart band may include data used previously in the individual smart bands but may also use new data from different sensors or data sources. In example embodiments, the expert system may be seeded with a plurality of smart bands and iterate to predict various states, but may also iterate towards reducing the number of smart bands needed to predict the same set of states.

[0937] Iteration of the expert system may be governed by rules, in some embodiments. For example, the expert system may be structured to collect data for seeding at a pre-determined frequency. The expert system may be structured to iterate at least a number of times, such as when a new component/equipment/fuel source is added, when a sensor goes off-line, or as standard practice. For example, when a sensor measuring the rotation of a stirrer in a food processing line goes off-line and the expert system begins acquiring data from a new sensor measuring the same data points, the expert system may be structured to iterate for a number of times before the state is utilized in or allowed to affect any downstream actions. The expert system may be structured to train off-line or train in situ/online. The expert system may be structured to include static and/or manually input data in its smart bands. For example, an expert system managing smart bands associated with a mixer in a food processing plant may be structured to iterate towards predicting a duration of mixing before the food being processed achieves a particular viscosity, wherein the smart band includes data regarding the speed of the mixer, temperature of its contents, viscometric measurements and the required endpoint for viscosity and temperature of the food. The expert system may be structured to include a minimum/maximum number of variables.

[0938] In example embodiments, the expert system may be overruled. In example embodiments, the expert system may revert to prior band settings, such as in the event the expert system fails, such as if a neural network fails in a neural net expert system, if uncertainty is too high in a model-based system, if the system is unable to resolve conflicting rules in rule-based system, or the system cannot converge on a solution in any of the foregoing. For example, sensor data on an irrigation system used by the expert system in a smart band may indicate a massive leak in the field, but visual inspection, such as by a drone, indicates no such leak. In this event, the expert system will revert to an original smart band for seeding the expert system. In another example, one or more point sensors on an industrial pressure cooker indicates imminent failure in a seal, but the data collection band that the expert system converged to with a weighting towards a performance metric did not identify the failure. In this event, the smart band will revert to an original setting or a version of the smart band that would have also identified the imminent failure of the pressure cooker seal. In example embodiments, the expert system may change smart band settings in the event that a new component is added that makes the system closer to a different offset system. For example, a vacuum distillation unit is added to an oil & gas refinery to distill naphthalene, but the current

smart band settings for the expert system are derived from a refinery that distills kerosene. In this example, a data structure with smart band settings for various offset systems may be searched for a system that is more closely matched to the current system. When a new offset system is identified as more closely matched, such as one that also distill naphthalene, the new smart band settings (e.g., which sensors to use, where to place them, how frequently to sample, what static data points are needed, etc. as described herein) are used to seed the expert system to iterate towards predicting a state for the system. In example embodiments, the expert system may change smart band settings in the event that a new set of offset data is available from a third-party library. For example, a pharmaceutical processing plant may have optimized a catalytic reactor to operate in a highly efficient way and deposited the smart band settings in a data structure. The data structure may be continuously scanned for new smart bands that better aid in monitoring catalytic reactions and thus, result in optimizing the operation of the reactor.

[0939] In example embodiments, the expert system may be used to uncover unknown variables. For example, the expert system may iterate to identify a missing variable to be used for further iterations, such as further neural net iterations. For example, an under-utilized tank in a legacy condensate/make-up water system of a power station may have an unknown capacity because it is inaccessible and no documentation exists on the tank. Various aspects of the tank may be measured by a swarm of sensors to arrive at an estimated volume (e.g., flow into a downstream space, duration of a dye traced solution to work through the system), then that volume can be fed into the neural net as a new variable in the smart band.

[0940] In example embodiments, the location of expert system node locations may be on a machine, on a data collector (or a group of them), in a network infrastructure (enterprise or other), or in the cloud. In example embodiments, there may be distributed neurons across nodes (e.g., machine, data collector, network, cloud).

[0941] In an aspect, a monitoring system 10700 for data collection in an industrial environment, comprising a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state. The data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state. The state may correspond to an outcome relating to a machine in the environment, an anticipated outcome relating to a machine in the environment, an outcome relating to a process in the environment, an anticipated outcome relating to a process in the environment, and the like. The collection parameter may be a bandwidth parameter, may be used to govern the multiplexing of a plurality of the input sensors, may be a timing parameter, may relate to a frequency range, may relate to the granularity of collection of sensor data, is a storage parameter for the collected data. The machine learning data analysis circuit may be structured to learn received output data patterns 10718 by being seeded with a model 10720, which may be a

physical model, an operational model, or a system model. The machine learning data analysis circuit may be structured to learn received output data patterns 10718 based on the state. The data collection band circuit may alter the subset of the plurality of sensors when the learned received output data pattern does not reliably predict the state, which may include discontinuing collection of data from the at least one subset.

[0942] The monitoring system 10700 may keep or modify operational parameters of an item of equipment in the environment based on the determined state. The controller 10706 may adjust the weighting of the machine learning data analysis circuit 10712 based on the learned received output data patterns 10718 or the state. The controller 10706 may collect more/fewer data points from one or more members of the at least one subset of plurality of sensors 10702 based on the learned received output data patterns 10718 or the state. The controller 10706 may change a data storage technique for the output data 10710 based on the learned received output data patterns 10718 or the state. The controller 10706 may change a data presentation mode or manner based on the learned received output data patterns 10718 or the state. The controller 10706 may apply one or more filters to the output data 10710. The controller 10706 may identify a new data collection band circuit 10708 based on one or more of the learned received output data patterns 10718 and the state. The controller 10706 may adjust the weights/biases of the machine learning data analysis circuit 10712, such as in response to the learned received output data patterns 10718, in response to the accuracy of the prediction of an anticipated state by the machine learning data analysis circuit, in response to the accuracy of a classification of a state by the machine learning data analysis circuit, and the like. The monitoring device 10700 may remove or re-task under-utilized equipment based on one or more of the learned received output data patterns 10718 and the state. The machine learning data analysis circuit 10712 may include a neural network expert system. At least one subset of the plurality of sensors measures vibration and noise data. The machine learning data analysis circuit 10712 may be structured to learn received output data patterns 10718 indicative of progress/alignment with one or more goals/guidelines, wherein progress/alignment of each goal/guideline may be determined by a different subset of the plurality of sensors. The machine learning data analysis circuit 10712 may be structured to learn received output data patterns 10718 indicative of an unknown variable. The machine learning data analysis circuit 10712 may be structured to learn received output data patterns 10718 indicative of a preferred input among available inputs. The machine learning data analysis circuit 10712 may be structured to learn received output data patterns 10718 indicative of a preferred input data collection band among available input data collection bands. The machine learning data analysis circuit 10712 may be disposed in part on a machine, on one or more data collectors, in network infrastructure, in the cloud, or any combination thereof.

[0943] In example embodiments, a monitoring device for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a controller 10706, the controller 10706 including a data collection band circuit 10708 structured to determine at least one subset of the plurality of sensors 10702 from which to process output data 10710; and a machine learning data analysis circuit 10712 structured to receive output data from

the at least one subset of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state, wherein the data collection band circuit 10708 alters an aspect of the at least one subset of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state. The aspect that the data collection band circuit 10708 alters is a number or a frequency of data points collected from one or more members of the at least one subset of plurality of sensors 10702. The aspect that the data collection band circuit 10708 alters is a bandwidth parameter, a timing parameter, a frequency range, a granularity of collection of sensor data, a storage parameter for the collected data, and the like.

[0944] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and wherein the data collection band circuit 10708 alters the at least one of the plurality of sensors 10702 when the learned received output data pattern 10718 does not reliably predict the state.

[0945] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and wherein the data collector 10704 collects more or fewer data points from the at least one of the plurality of sensors 10702 based on the learned received output data patterns 10718 or the state.

[0946] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data 10710 patterns indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and

wherein the controller 10706 changes a data storage technique for the output data 10710 based on the learned received output data patterns 10718 or the state.

[0947] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and wherein the controller 10706 changes a data presentation mode or manner based on the learned received output data patterns 10718 or the state.

[0948] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and wherein the controller 10706 identifies a new data collection band circuit 10708 based on one or more of the learned received output data patterns 10718 and the state.

[0949] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712 structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and wherein the controller 10706 adjusts the weights/biases of the machine learning data analysis circuit 10712. The adjustment may be in response to the learned received output data patterns, in response to the accuracy of the prediction of an anticipated state by the machine learning data analysis circuit, in response to the accuracy of a classification of a state by the machine learning data analysis circuit, and the like.

[0950] In example embodiments, a monitoring system 10700 for data collection in an industrial environment may include a plurality of input sensors 10702 communicatively coupled to a data

collector 10704 having a controller 10706, a data collection band circuit 10708 structured to determine at least one collection parameter for at least one of the plurality of sensors 10702 from which to process output data 10710, and a machine learning data analysis circuit 10712. This machine learning data analysis circuit is structured to receive output data 10710 from the at least one of the plurality of sensors 10702 and learn received output data patterns 10718 indicative of a state, wherein the data collection band circuit 10708 alters the at least one collection parameter for the at least one of the plurality of sensors 10702 based on one or more of the learned received output data patterns 10718 and the state, and wherein the machine learning data analysis circuit 10712 is structured to learn received output data patterns 10718 indicative of progress or alignment with one or more goals or guidelines.

[0951] As described elsewhere herein, an expert system in an industrial environment may use sensor data to make predictions about outcomes or states of the environment or items in the environment. Data collection may be of various types of data (e.g., vibration data, noise data and other sensor data of the types described throughout this disclosure) for event detection, state detection, and the like. For example, the expert system may utilize ambient noise, or the overall sound environment of the area and/or overall vibration of the device of interest, optionally in conjunction with other sensor data, in detecting or predicting events or states. For example, a reciprocating compressor in a refinery, which may generate its own vibration, may also have an ambient vibration through contact with other aspects of the system.

[0952] In example embodiments, all three types of noise (ambient noise, local noise and vibration noise) including various subsets thereof and combinations with other types of data, may be organized into large data sets, along with measured results, that are processed by a "deep learning" machine/expert system that learns to predict one or more states (e.g., maintenance, failure, or operational) or overall outcomes, such as by learning from human supervision or from other feedback, such as feedback from one or more of the systems described throughout this disclosure and the documents incorporated by reference herein.

[0953] Throughout this disclosure, various examples will involve machines, components, equipment, assemblies, and the like, and it should be understood that the disclosure could apply to any of the aforementioned. Elements of these machines operating in an industrial environment (e.g., rotating elements, reciprocating elements, swinging elements, flexing elements, flowing elements, suspending elements, floating elements, bouncing elements, bearing elements, etc.) may generate vibrations that may be of a specific frequency and/or amplitude typical of the element when the element is in a given operating condition or state (e.g., a normal mode of operation of a machine at a given speed, in a given gear, or the like). Changes in a parameter of the vibration may be indicative or predictive of a state or outcome of the machine. Various sensors may be useful in measuring vibration, such as accelerometers, velocity transducers, imaging sensors, acoustic sensors, and displacement probes, which may collectively be known as vibration sensors. Vibration sensors may be mounted to the machine, such as permanently or temporarily (e.g., adhesive, hook-and-loop, or magnetic attachment), or may be disposed on a mobile or portable data collector. Sensed conditions may be compared to historical data to identify or predict a state, condition or

outcome. Typical faults that can be identified using vibration analysis include: machine out of balance, machine out of alignment, resonance, bent shafts, gear mesh disturbances, blade pass disturbances, vane pass disturbances, recirculation & cavitation, motor faults (rotor & stator), bearing failures, mechanical looseness, critical machine speeds, and the like, as well as excessive friction, clutch slipping, belt problems, suspension and shock absorption problems, valve and other fluid leaks, under-pressure states in lubrication and other fluid systems, overheating (such as due to many of the above), blockage or freezing of engagement of mechanical systems, interference effects, and other faults described throughout this disclosure and in the documents incorporated by reference.

[0954] Given that machines are frequently found adjacent to or working in concert with other machinery, measuring the vibration of the machine may be complicated by the presence of various noise components in the environment or associated vibrations that the machine may be subjected to. Indeed, the ambient and/or local environment may have its own vibration and/or noise pattern that may be known. In example embodiments, the combination of vibration data with ambient and/or local noise or other ambient sensed conditions may form its own pattern, as will be further described herein.

[0955] In example embodiments, measuring vibration noise may involve one or more vibration sensors on or in a machine to measure vibration noise of the machine that occurs continuously or periodically. Analysis of the vibration noise may be performed, such as filtering, signal conditioning, spectral analysis, trend analysis, and the like. Analysis may be performed on aggregate or individual sensor measurements to isolate vibration noise of equipment to obtain a characteristic vibration, vibration pattern or "vibration fingerprint" of the machine. The vibration fingerprints may be stored in a data structure, or library, of vibration fingerprints. The vibration fingerprints may include frequencies, spectra (i.e., frequency vs. amplitude), velocities, peak locations, wave peak shapes, waveform shapes, wave envelope shapes, accelerations, phase information, phase shifts (including complex phase measurements) and the like. Vibration fingerprints may be stored in the library in association with a parameter by which it may be searched or sorted. The parameters may include a brand or type of machine/component/equipment, location of sensor(s) attachment or placement, duty cycle of the equipment/machine, load sharing of the equipment/machine, dynamic interactions with other devices, RPM, flow rate, pressure, other vibration driving characteristic, voltage of line power, age of equipment, time of operation, known neighboring equipment, associated auxiliary equipment/components, size of space equipment is in, material of platform for equipment, heat flux, magnetic fields, electrical fields, currents, voltage, capacitance, inductance, aspect of a product, and combinations (e.g., simple ratios) of the same. Vibration fingerprints may be obtained for machines under normal operation or for other periods of operation (e.g., off-nominal operation, malfunction, maintenance needed, faulty component, incorrect parameters of operation, other conditions, etc.) and can be stored in the library for comparison to current data. The library of vibration fingerprints may be stored as indicators with associated predictions, states, outcomes and/or events. Trend analysis data of measured vibration fingerprints can indicate time between maintenance events/failure events.

[0956] In example embodiments, vibration noise may be used by the expert system to confirm the status of a machine, such as a favorable operation, a production rate, a generation rate, an operational efficiency, a financial efficiency (e.g., output per cost), a power efficiency, and the like. In example embodiments, the expert system may make a comparison of the vibration noise with a stored vibration fingerprint. In other embodiments, the expert system may be seeded with vibration noise and initial feedback on states and outcomes in order to learn to predict other states and outcomes. For example, a center pivot irrigation system may be remotely monitored by attached vibration sensors to provide a measured vibration noise that can be compared to a library of vibration fingerprints to confirm that the system is operating normally. If the system is not operating normally, the expert system may automatically dispatch a field crew or drone to investigate. In another example of a vacuum distillation unit in a refinery, the vibration noise may be compared, such as by the expert system, to stored vibration fingerprints in a library to confirm a production rate of diesel. In a further example, the expert system may be seeded with vibration noise for a pipeline under conditions of a normal production rate and as the expert system iterates with current data (e.g., altered vibration noise, and possibly other altered parameters), it may predict that the production rate has increased as caused by the alterations. Measurements may be continually analyzed in this way to remotely monitor operation.

[0957] In example embodiments, vibration noise may be compared, such as by the expert system, to stored vibration fingerprints and associated states and outcomes in the library, or alternatively, may be used to seed an expert system to predict when maintenance is required (e.g., off-nominal measurement, artifacts in signal, etc.), such as when vibration noise is matched to a condition when the equipment/component required maintenance, vibration noise exceeds a threshold/limit, vibration noise exceeds a threshold/limit or matches a library vibration fingerprint together with one or more additional parameters, as described herein. For example, when the vibration fingerprint from a turbine agitator in a pharmaceutical processing plant matches a vibration fingerprint for a turbine agitator when it required a replacement bearing, the expert system may cause an action to occur, such as immediately shutting down the agitator or scheduling its shutdown and maintenance.

[0958] In example embodiments, vibration noise may be compared, such as by the expert system, to stored vibration fingerprints and associated states and outcomes in the library, or alternatively, may be used to seed an expert system to predict a failure or an imminent failure. For example, vibration noise from a gas agitator in a pharmaceutical processing plant may be matched to a condition when the agitator previously failed or was about to fail. In this example, the expert system may immediately shut down the agitator, schedule its shutdown, or cause a backup agitator to come online. In another example, vibration noise from a pump blasting liquid agitator in a chemical processing plant may exceed a threshold or limit and the expert system may cause an investigation into the cause of the excess vibration noise, shut down the agitator, or the like. In another example, vibration noise from an anchor agitator in a pharmaceutical processing plant may exceed a threshold/limit or match a library vibration fingerprint together with one or more additional parameters (see parameters herein), such as a decreased flow rate, increased temperature, or the

like. Using vibration noise taken together with the parameters, the expert system may more reliably predict the failure or imminent failure.

[0959] In example embodiments, vibration noise may be compared, such as by the expert system, to stored vibration fingerprints and associated states and outcomes in the library, or alternatively, may be used to seed an expert system to predict or diagnose a problem (e.g., unbalanced, misaligned, worn, or damaged) with the equipment or an external source contributing vibration noise to the equipment. For example, when the vibration noise from a paddle-type agitator mixer matches a vibration fingerprint from a prior imbalance, the expert system may immediately shut down the mixer.

[0960] In example embodiments, when the expert system makes a prediction of an outcome or state using vibration noise, the expert system may perform a downstream action, or cause it to be performed. Downstream actions may include: triggering an alert of a failure, imminent failure, or maintenance event; shutting down equipment/component; initiating maintenance/lubrication/alignment; deploying a field technician; recommending a vibration absorption/dampening device; modifying a process to utilize backup equipment/component; modifying a process to preserve products/reactants, etc.; generating/modifying a maintenance schedule; coupling the vibration fingerprint with duty cycle of the equipment, RPM, flow rate, pressure, temperature or other vibration-driving characteristic to obtain equipment/component status and generate a report, and the like. For example, vibration noise for a catalytic reactor in a chemical processing plant may be matched to a condition when the catalytic reactor required maintenance. Based on this predicted state of required maintenance, the expert system may deploy a field technician to perform the maintenance.

[0961] In example embodiments, the library may be updated if a changed parameter resulted in a new vibration fingerprint, or if a predicted outcome or state did not occur in the absence of mitigation. In example embodiments, the library may be updated if a vibration fingerprint was associated with an alternative state than what was predicted by the library. The update may occur after just one time that the state that actually occurred did not match the predicted state from the library. In other embodiments, it may occur after a threshold number of times. In example embodiments, the library may be updated to apply one or more rules for comparison, such as rules that govern how many parameters to match along with the vibration fingerprint, or the standard deviation for the match in order to accept the predicted outcome.

[0962] In example embodiments, vibration noise may be compared, such as by the expert system, to stored vibration fingerprints and associated states and outcomes in the library, or alternatively, may be used to seed an expert system to determine if a change in a system parameter external or internal to the machine has an effect on its intrinsic operation. In example embodiments, a change in one or more of a temperature, flow rate, materials in use, duration of use, power source, installation, or other parameter (see parameters above) may alter the vibration fingerprint of a machine. For example, in a pressure reactor in a chemical processing plant, the flow rate and a reactant may be changed. The changes may alter the vibration fingerprint of the machine such that the vibration fingerprint stored in the library for normal operation is no longer correct.

[0963] Ambient noise, or the overall sound environment of the area and/or overall vibration of the device of interest, optionally in conjunction with other ambient sensed conditions, may be used in detecting or predicting events, outcomes, or states. Ambient noise may be measured by a microphone, ultrasound sensors, acoustic wave sensors, optical vibration sensors (e.g., using a camera to see oscillations that produce noise), or "deep learning" neural networks involving various sensor arrays that learn, using large data sets, to identify patterns, sounds types, noise types, etc. In example embodiments, the ambient sensed condition may relate to motion detection. For example, the motion may be a platform motion (e.g., vehicle, oil platform, suspended platform on land, etc.) or an object motion (e.g., moving equipment, people, robots, parts (e.g., fan blades or turbine blades), etc.). In example embodiments, the ambient sensed condition may be sensed by imaging, such as to detect a location and nature of various machines, equipment, and other objects, such as ones that might impact local vibration. In example embodiments, the ambient sensed condition may be sensed by thermal detection and imaging (e.g., for presence of people; presence of heat sources that may affect performance parameters, etc.). In example embodiments, the ambient sensed condition may be sensed by field detection (e.g., electrical, magnetic, etc.). In example embodiments, the ambient sensed condition may be sensed by chemical detection (e.g., smoke, other conditions). Any sensor data may be used by the expert system to provide an ambient sensed condition for analysis along with the vibration fingerprint to predict an outcome, event, or state. For example, an ambient sensed condition near a stirrer or mixer in a food processing plant may be the operation of a space heater during winter months, wherein the ambient sensed condition may include an ambient noise and an ambient temperature.

[0964] In an aspect, local noise may be the noise or vibration environment, which is ambient, but known to be locally generated. The expert system may filter out ambient noise, employ common mode noise removal, and/or physically isolate the sensing environment.

[0965] In example embodiments, a system for data collection in an industrial environment may use ambient, local and vibration noise for prediction of outcomes, events, and states. A library may be populated with each of the three noise types for various conditions (e.g., start up, shut down, normal operation, other periods of operation as described elsewhere herein). In other embodiments, the library may be populated with noise patterns representing the aggregate ambient, local, and/or vibration noise. Analysis (e.g., filtering, signal conditioning, spectral analysis, trend analysis) may be performed on the aggregate noise to obtain a characteristic noise pattern and identify changes in noise pattern as possible indicators of a changed condition. A library of noise patterns may be generated with established vibration fingerprints and local and ambient noise that can be sorted by a parameter (see parameters herein), or other parameters/features of the local and ambient environment (e.g., company type, industry type, products, robotic handling unit present/not present, operating environment, flow rates, production rates, brand or type of auxiliary equipment (e.g., filters, seals, coupled machinery)). The library of noise patterns may be used by an expert system, such as one with machine learning capacity, to confirm a status of a machine, predict when maintenance is required (e.g., off-nominal measurement, artifacts in signal), predict a failure or an imminent failure, predict/diagnose a problem, and the like.

[0966] Based on a current noise pattern, the library may be consulted or used to seed an expert system to predict an outcome, event, or state based on the noise pattern. Based on the prediction, the expert system may one or more of trigger an alert of a failure, imminent failure, or maintenance event, shut down equipment/component/line, initiate maintenance/lubrication/alignment, deploy a field technician, recommend a vibration absorption/dampening device, modify a process to utilize backup equipment/component, modify a process to preserve products/reactants, etc., generate/modify a maintenance schedule, or the like.

[0967] For example, a noise pattern for a thermic heating system in a pharmaceutical plant or cooking system may include local, ambient, and vibration noise. The ambient noise may be a result of, for example, various pumps to pump fuel into the system. Local noise may be a result of a local security camera chirping with every detection of motion. Vibration noise may result from the combustion machinery used to heat the thermal fluid. These noise sources may form a noise pattern which may be associated with a state of the thermic system. The noise pattern and associated state may be stored in a library. An expert system used to monitor the state of the thermic heating system may be seeded with noise patterns and associated states from the library. As current data are received into the expert system, it may predict a state based on having learned noise patterns and associated states.

[0968] In another example, a noise pattern for boiler feed water in a refinery may include local and ambient noise. The local noise may be attributed to the operation of, for example, a feed pump feeding the feed water into a steam drum. The ambient noise may be attributed to nearby fans. These noise sources may form a noise pattern which may be associated with a state of the boiler feed water. The noise pattern and associated state may be stored in a library. An expert system used to monitor the state of the boiler may be seeded with noise patterns and associated states from the library. As current data are received into the expert system, it may predict a state based on having learned noise patterns and associated states.

[0969] In yet another example, a noise pattern for a storage tank in a refinery may include local, ambient, and vibration noise. The ambient noise may be a result of, for example, a pump that pumps a product into the tank. Local noise may be a result of a fan ventilating the tank room. Vibration noise may result from line noise of a power supply into the storage tank. These noise sources may form a noise pattern which may be associated with a state of the storage tank. The noise pattern and associated state may be stored in a library. An expert system used to monitor the state of the storage tank may be seeded with noise patterns and associated states from the library. As current data are received into the expert system, it may predict a state based on having learned noise patterns and associated states.

[0970] In another example, a noise pattern for condensate/make-up water system in a power station may include vibration and ambient noise. The ambient noise may be attributed to nearby fans. The vibration noise may be attributed to the operation of the condenser. These noise sources may form a noise pattern which may be associated with a state of the condensate/make-up water system. The noise pattern and associated state may be stored in a library. An expert system used to monitor the state of the condensate/make-up water system may be seeded with noise patterns and associated

states from the library. As current data are received into the expert system, it may predict a state based on having learned noise patterns and associated states.

[0971] A library of noise patterns may be updated if a changed parameter resulted in a new noise pattern or if a predicted outcome or state did not occur in the absence of mitigation of a diagnosed problem. A library of noise patterns may be updated if a noise pattern resulted in an alternative state than what was predicted by the library. The update may occur after just one time that the state that actually occurred did not match the predicted state from the library. In other embodiments, it may occur after a threshold number of times. In example embodiments, the library may be updated to apply one or more rules for comparison, such as rules that govern how many parameters to match along with the noise pattern, or the standard deviation for the match in order to accept the predicted outcome. For example, a baffle may be replaced in a static agitator in a pharmaceutical processing plant which may result in a changed noise pattern. In another example, as the seal on a pressure cooker in a food processing plant ages, the noise pattern associated with the pressure cooker may change.

[0972] In example embodiments, the library of vibration fingerprints, noise sources and/or noise patterns may be available for subscription. The libraries may be used in offset systems to improve operation of the local system. Subscribers may subscribe at any level (e.g., component, machinery, installation, etc.) in order to access data that would normally not be available to them, such as because it is from a competitor, or is from an installation of the machinery in a different industry not typically considered. Subscribers may search on indicators/predictors based on or filtered by system conditions, or update an indicator/predictor with proprietary data to customize the library. The library may further include parameters and metadata auto-generated by deployed sensors throughout an installation, onboard diagnostic systems and instrumentation and sensors, ambient sensors in the environment, sensors (e.g., in flexible sets) that can be put into place temporarily, such as in one or more mobile data collectors, sensors that can be put into place for longer term use, such as being attached to points of interest on devices or systems, and the like.

[0973] In example embodiments, a third party (e.g., RMOs, manufacturers) can aggregate data at the component level, equipment level, factory/installation level and provide a statistically valid data set against which to optimize their own systems. For example, when a new installation of a machine is contemplated, it may be beneficial to review a library for the best data points to acquire in making state predictions. For example, a particular sensor package may be recommended to reliably determine if there will be a failure. For example, if vibration noise of equipment coupled with particular levels of local noise or other ambient sensed conditions reliably is an indicator of imminent failure, a given vibration transducer/temp/microphone package observing those elements may be recommended for the installation. Knowing such information may inform the choice to rent or buy a piece of machinery or associated warranties and service plans, such as based on knowing the quantity and depth of information that may be needed to reliably maintain the machinery.

[0974] In example embodiments, manufacturers may utilize the library to rapidly collect in-service information for machines to draft engineering specifications for new customers.

[0975] In example embodiments, noise and vibration data may be used to remotely monitor installs and automatically dispatch a field crew.

[0976] In example embodiments, noise and vibration data may be used to audit a system. For example, equipment running outside the range of a licensed duty cycle may be detected by a suite of vibration sensors and/or ambient/local noise sensors. In example embodiments, alerts may be triggered of potential out-of-warranty violations based on data from vibration sensors and/or ambient/local noise sensors.

[0977] In example embodiments, noise and vibration data may be used in maintenance. This may be particularly useful where multiple machines are deployed that may vibrationally interact with the environment, such as two large generating machines on the same floor or platform with each other, such as in power generation plants.

[0978] In example embodiments, a monitoring system 10800 for data collection in an industrial environment, may include a plurality of sensors 10802 selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to a data collector 10804, a data collection circuit 10808 structured to collect output data 10810 from the plurality of sensors 10802, and a machine learning data analysis circuit 10812 structured to receive the output data 10810 and learn received output data patterns 10814 predictive of at least one of an outcome and a state. The state may correspond to an outcome relating to a machine in the environment, an anticipated outcome relating to a machine in the environment, an outcome relating to a process in the environment, or an anticipated outcome relating to a process in the environment. The system may be deployed on the data collector 10804 or distributed between the data collector 10804 and a remote infrastructure. The data collector 10804 may include the data collection circuit 10808. The ambient environment condition or local sensors include one or more of a noise sensor, a temperature sensor, a flow sensor, a pressure sensor, a chemical sensor, a vibration sensor, an acceleration sensor, an accelerometer, a Pressure sensor, a force sensor, a position sensor, a location sensor, a velocity sensor, a displacement sensor, a temperature sensor, a thermographic sensor, a heat flux sensor, a tachometer sensor, a motion sensor, a magnetic field sensor, an electrical field sensor, a galvanic sensor, a current sensor, a flow sensor, a gaseous flow sensor, a non-gaseous fluid flow sensor, a heat flow sensor, a particulate flow sensor, a level sensor, a proximity sensor, a toxic gas sensor, a chemical sensor, a CBRNE sensor, a pH sensor, a hygrometer, a moisture sensor, a densitometer, an imaging sensor, a camera, an SSR, a triax probe, an ultrasonic sensor, a touch sensor, a microphone, a capacitive sensor, a strain gauge, an EMF meter, and the like.

[0979] In example embodiments, a monitoring system 10800 for data collection in an industrial environment may include a data collection circuit 10808 structured to collect output data 10810 from a plurality of sensors 10802 selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to a data collection circuit 10808, and a machine learning data analysis circuit 10812 structured to receive the output data

10810 and learn received output data patterns 10814 predictive of at least one of an outcome and a state, wherein the monitoring system 10800 is structured to determine if the output data matches a learned received output data pattern. The machine learning data analysis circuit 10812 may be structured to learn received output data patterns 10814 by being seeded with a model 10816. The model 10816 may be a physical model, an operational model, or a system model. The machine learning data analysis circuit 10812 may be structured to learn received output data patterns 10814 based on the outcome or the state. The monitoring system 10700 keeps or modifies operational parameters or equipment based on the predicted outcome or the state. The data collection circuit 10808 collects more or fewer data points from one or more of the plurality of sensors 10802 based on the learned received output data patterns 10814, the outcome or the state. The data collection circuit 10808 changes a data storage technique for the output data based on the learned received output data patterns 10814, the outcome, or the state. The data collector 10804 changes a data presentation mode or manner based on the learned received output data patterns 10814, the outcome, or the state. The data collection circuit 10808 applies one or more filters (low pass, high pass, band pass, etc.) to the output data. The data collection circuit 10808 adjusts the weights/biases of the machine learning data analysis circuit 10812, such as in response to the learned received output data patterns 10814. The monitoring system 10800 removes/re-tasks under-utilized equipment based on one or more of the learned received output data patterns 10814, the outcome, or the state. The machine learning data analysis circuit 10812 may include a neural network expert system. The machine learning data analysis circuit 10812 may be structured to learn received output data patterns 10814 indicative of progress/alignment with one or more goals/guidelines, wherein progress/alignment of each goal/guideline is determined by a different subset of the plurality of sensors 10802. The machine learning data analysis circuit 10812 may be structured to learn received output data patterns 10814 indicative of an unknown variable. The machine learning data analysis circuit 10812 may be structured to learn received output data patterns 10814 indicative of a preferred input sensor among available input sensors. The machine learning data analysis circuit 10812 may be disposed in part on a machine, on one or more data collection circuits 10808, in network infrastructure, in the cloud, or any combination thereof. The output data 10810 from the vibration sensors forms a vibration fingerprint, which may include one or more of a frequency, a spectrum, a velocity, a peak location, a wave peak shape, a waveform shape, a wave envelope shape, an acceleration, a phase information, and a phase shift. The data collection circuit 10808 may apply a rule regarding how many parameters of the vibration fingerprint to match or the standard deviation for the match in order to identify a match between the output data 10810 and the learned received output data pattern. The state may be one of a normal operation, a maintenance required, a failure, or an imminent failure. The monitoring system 10800 may trigger an alert, shut down equipment/component/line, initiate maintenance/lubrication/alignment based on the predicted outcome or state, deploy a field technician based on the predicted outcome or state, recommend a vibration absorption/dampening device based on the predicted outcome or state, modify a process to utilize backup equipment/component based on the predicted outcome or state, and the like. The monitoring system 10800 may modify a process to preserve

products/reactants, etc. based on the predicted outcome or state. The monitoring system 10800 may generate or modify a maintenance schedule based on the predicted outcome or state. The data collection circuit 10808 may include the data collection circuit 10808. The system may be deployed on the data collection circuit 10808 or distributed between the data collection circuit 10808 and a remote infrastructure.

[0980] In example embodiments, a monitoring system 10800 for data collection in an industrial environment may include a data collection circuit 10808 structured to collect output data 10810 from a plurality of sensors 10802 selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to the data collection circuit 10808, and a machine learning data analysis circuit 10812 structured to receive the output data 10810 and learn received output data patterns 10814 predictive of at least one of an outcome and a state, wherein the monitoring system 10800 is structured to determine if the output data matches a learned received output data pattern and keep or modify operational parameters or equipment based on the determination.

[0981] In example embodiments, a monitoring system 10800 for data collection in an industrial environment may include a data collection circuit 10808 structured to collect output data 10810 from the plurality of sensors 10802 selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to the data collection circuit 10808, and a machine learning data analysis circuit 10812 structured to receive the output data 10810 and learn received output data patterns 10814 predictive of at least one of an outcome and a state, wherein the output data 10810 from the vibration sensors forms a vibration fingerprint. The vibration fingerprint may include one or more of a frequency, a spectrum, a velocity, a peak location, a wave peak shape, a waveform shape, a wave envelope shape, an acceleration, a phase information, and a phase shift. The data collection circuit 10808 may apply a rule regarding how many parameters of the vibration fingerprint to match or the standard deviation for the match in order to identify a match between the output data 10810 and the learned received output data pattern. The monitoring system 10800 may be structured to determine if the output data matches a learned received output data pattern and keep or modify operational parameters or equipment based on the determination.

[0982] In example embodiments, a monitoring system 10800 for data collection in an industrial environment may include a data collection band circuit 10818 that identifies a subset of the plurality of sensors 10802 from which to process output data, the sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to a data collection band circuit 10818, a data collection circuit 10808 structured to collect the output data 10810 from the subset of plurality of sensors 10802, and a machine learning data analysis circuit 10812 structured to receive the output data 10810 and learn received output data patterns 10814 predictive of at least one of an outcome and a state, wherein when the learned

received output data patterns 10814 do not reliably predict the outcome or the state, the data collection band circuit 10818 alters at least one parameter of at least one of the plurality of sensors 10802. A controller 10806 identifies a new data collection band circuit 10818 based on one or more of the learned received output data patterns 10814 and the outcome or state. The machine learning data analysis circuit 10812 may be further structured to learn received output data patterns 10814 indicative of a preferred input data collection band among available input data collection bands. The system may be deployed on the data collection circuit 10808 or distributed between the data collection circuit 10808 and a remote infrastructure.

[0983] In example embodiments, a monitoring system for data collection in an industrial environment may include a data collection circuit 10808 structured to collect output data 10810 from a plurality of sensors 10802, the sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to the data collection circuit 10808, wherein the output data 10810 from the vibration sensors is in the form of a vibration fingerprint, a data structure 10820 comprising a plurality of vibration fingerprints and associated outcomes, and a machine learning data analysis circuit 10812 structured to receive the output data 10810 and learn received output data patterns 10814 predictive of an outcome or a state based on processing of the vibration fingerprints. The machine learning data analysis circuit 10812 may be seeded with one of the plurality of vibration fingerprints from the data structure 10820. The data structure 10820 may be updated if a changed parameter resulted in a new vibration fingerprint or if a predicted outcome did not occur in the absence of mitigation. The data structure 10820 may be updated when the learned received output data patterns 10814 do not reliably predict the outcome or the state. The system may be deployed on the data collection circuit or distributed between the data collection circuit and a remote infrastructure.

[0984] In example embodiments, a monitoring system 10800 for data collection in an industrial environment may include a data collection circuit 10808 structured to collect output data 10810 from a plurality of sensors 10802 selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors 10802 communicatively coupled to a data collection circuit 10808, wherein the output data 10810 from the plurality of sensors 10802 is in the form of a noise pattern, a data structure 10820 comprising a plurality of noise patterns and associated outcomes, and a machine learning data analysis circuit 10812 structured to receive the output data 10810 and learn received output data patterns 10814 predictive of an outcome or a state based on processing of the noise patterns.

[0985] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a plurality of sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors communicatively coupled to a data collector; a data collection circuit structured to collect output data from the plurality of sensors; and a machine learning data analysis circuit structured to receive the output data and learn received

output data patterns predictive of at least one of an outcome and a state. The state may correspond to an outcome, anticipated outcome, outcome relating to a process, as relating to a machine in the environment. The system may be deployed on the data collector. The system may be distributed between the data collector and a remote infrastructure. The ambient environment condition sensors may include a noise sensor, a temperature sensor, a flow sensor, a pressure sensor, include a chemical sensor, a noise sensor, a temperature sensor, a flow sensor, a pressure sensor, a chemical sensor, a vibration sensor, an acceleration sensor, an accelerometer, a pressure sensor, a force sensor, a position sensor, a location sensor, a velocity sensor, a displacement sensor, a temperature sensor, a thermographic sensor, a heat flux sensor, a tachometer sensor, a motion sensor, a magnetic field sensor, an electrical field sensor, a galvanic sensor, a current sensor, a flow sensor, a gaseous flow sensor, a non-gaseous fluid flow sensor, a heat flow sensor, a particulate flow sensor, a level sensor, a proximity sensor, a toxic gas sensor, a chemical sensor, a CBRNE sensor, a pH sensor, a hygrometer, a moisture sensor, a densitometer, an imaging sensor, a camera, an SSR, a triax probe, an ultrasonic sensor, a touch sensor, a microphone, a capacitive sensor, a strain gauge, and an EMF meter. The local sensors may comprise one or more of a vibration sensor, an acceleration sensor, an accelerometer, a pressure sensor, a force sensor, a position sensor, a location sensor, a velocity sensor, a displacement sensor, a temperature sensor, a thermographic sensor, a heat flux sensor, a tachometer sensor, a motion sensor, a magnetic field sensor, an electrical field sensor, a galvanic sensor, a current sensor, a flow sensor, a gaseous flow sensor, a non-gaseous fluid flow sensor, a heat flow sensor, a particulate flow sensor, a level sensor, a proximity sensor, a toxic gas sensor, a chemical sensor, a CBRNE sensor, a pH sensor, a hygrometer, a moisture sensor, a densitometer, an imaging sensor, a camera, an SSR, a triax probe, an ultrasonic sensor, a touch sensor, a microphone, a capacitive sensor, a strain gauge, and an EMF meter.

[0986] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collection circuit structured to collect output data from a plurality of sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors communicatively coupled to the data collection circuit; and a machine learning data analysis circuit structured to receive the output data and learn received output data patterns predictive of at least one of an outcome and a state, wherein the monitoring system is structured to determine if the output data matches a learned received output data pattern. In example embodiments, the machine learning data analysis circuit may be structured to learn received output data patterns by being seeded with a model, such as where the model is a physical model, an operational model, or a system model. The machine learning data analysis circuit may be structured to learn received output data patterns based on the outcome or the state. The monitoring system may keep or modify operational parameters or equipment based on the predicted outcome or the state. The data collection circuit collects data points from one or more of the plurality of sensors based on the learned received output data patterns, the outcome, or the state. The data collection circuit may change a data storage technique for the output data based on the learned received output data patterns, the outcome, or the state. The data collection circuit may change a data presentation

mode or manner based on the learned received output data patterns, the outcome, or the state. The data collection circuit may apply one or more filters (low pass, high pass, band pass, etc.) to the output data. The data collection circuit may adjust the weights/biases of the machine learning data analysis circuit, such as where the adjustment is in response to the learned received output data patterns. The monitoring system may remove, or re-task under-utilized equipment based on one or more of the learned received output data patterns, the outcome, or the state. The machine learning data analysis circuit may include a neural network expert system. The machine learning data analysis circuit may be structured to learn received output data patterns indicative of progress/alignment with one or more goals or guidelines, such as where progress or alignment of each goal or guideline is determined by a different subset of the plurality of sensors. The machine learning data analysis circuit may be structured to learn received output data patterns indicative of an unknown variable. The machine learning data analysis circuit may be structured to learn received output data patterns indicative of a preferred input sensor among available input sensors. The machine learning data analysis circuit may be disposed in part on a machine, on one or more data collectors, in network infrastructure, in the cloud, or any combination thereof. The output data from the vibration sensors may form a vibration fingerprint, such as where the vibration fingerprint includes one or more of a frequency, a spectrum, a velocity, a peak location, a wave peak shape, a waveform shape, a wave envelope shape, an acceleration, a phase information, and a phase shift. The data collection circuit may apply a rule regarding how many parameters of the vibration fingerprint to match or the standard deviation for the match in order to identify a match between the output data and the learned received output data pattern. The state may be one of a normal operation, a maintenance required, a failure, or an imminent failure. The monitoring system may trigger an alert based on the predicted outcome or state. The monitoring system may shut down equipment, component, or line based on the predicted outcome or state. The monitoring system may initiate maintenance, lubrication, or alignment based on the predicted outcome or state. The monitoring system may deploy a field technician based on the predicted outcome or state. The monitoring system may recommend a vibration absorption or dampening device based on the predicted outcome or state. The monitoring system may modify a process to utilize backup equipment or a component based on the predicted outcome or state. The monitoring system may modify a process to preserve products or reactants based on the predicted outcome or state. The monitoring system may generate or modify a maintenance schedule based on the predicted outcome or state. The system may be distributed between the data collector and a remote infrastructure.

[0987] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collection circuit structured to collect output data from a plurality of sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors communicatively coupled to the data collection circuit; and a machine learning data analysis circuit structured to receive the output data and learn received output data patterns predictive of at least one of an outcome and a state, wherein the monitoring system is structured to

determine if the output data matches a learned received output data pattern and keep or modify operational parameters or equipment based on the determination.

[0988] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collection circuit structured to collect output data from a plurality of sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors communicatively coupled to the data collection circuit; and a machine learning data analysis circuit structured to receive the output data and learn received output data patterns predictive of at least one of an outcome and a state, wherein the output data from the vibration sensors forms a vibration fingerprint. In example embodiments, the vibration fingerprint may comprise one or more of a frequency, a spectrum, a velocity, a peak location, a wave peak shape, a waveform shape, a wave envelope shape, an acceleration, a phase information, and a phase shift. The data collection circuit may apply a rule regarding how many parameters of the vibration fingerprint to match or the standard deviation for the match in order to identify a match between the output data and the learned received output data pattern. The monitoring system may be structured to determine if the output data matches a learned received output data pattern and keep or modify operational parameters or equipment based on the determination.

[0989] In example embodiments, a monitoring system for data collection in an industrial environment may comprise: a data collection band circuit that identifies a subset of a plurality of sensors from which to process output data, the sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors communicatively coupled to the data collection band circuit; a data collection circuit structured to collect the output data from the subset of plurality of sensors; and a machine learning data analysis circuit structured to receive the output data and learn received output data patterns predictive of at least one of an outcome and a state wherein when the learned received output data patterns do not reliably predict the outcome or the state, the data collection band circuit alters at least one parameter of at least one of the plurality of sensors. In example embodiments, the controller may identify a new data collection band circuit based on one or more of the learned received output data patterns and the outcome or state. The machine learning data analysis circuit may be further structured to learn received output data patterns indicative of a preferred input data collection band among available input data collection bands. The system may be distributed between the data collection circuit and a remote infrastructure.

[0990] In example embodiments, a monitoring system for data collection in an industrial environment may comprise a data collection circuit structured to collect output data from the plurality of sensors, the sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment and being communicatively coupled to the data collection circuit, wherein the output data from the vibration sensors is in the form of a vibration fingerprint; a data structure comprising a plurality of vibration fingerprints and associated outcomes; and a machine learning data analysis

circuit structured to receive the output data and learn received output data patterns predictive of an outcome or a state based on processing of the vibration fingerprints. The machine learning data analysis circuit may be seeded with one of the plurality of vibration fingerprints from the data structure. The data structure may be updated if a changed parameter resulted in a new vibration fingerprint or if a predicted outcome did not occur in the absence of mitigation. The data structure may be updated when the learned received output data patterns do not reliably predict the outcome or the state. The system may be distributed between the data collection circuit and a remote infrastructure.

[0991] In example embodiments, a monitoring system for data collection in an industrial environment may comprise a data collection circuit structured to collect output data from the plurality of sensors selected among vibration sensors, ambient environment condition sensors and local sensors for collecting non-vibration data proximal to a machine in the environment, the plurality of sensors communicatively coupled to the data collection circuit, wherein the output data from the plurality of sensors is in the form of a noise pattern; a data structure comprising a plurality of noise patterns and associated outcomes; and a machine learning data analysis circuit structured to receive the output data and learn received output data patterns predictive of an outcome or a state based on processing of the noise patterns.

[0992] An example system for data collection in an industrial environment includes an industrial system having a number of components, and a number of sensors wherein each of the sensors is operatively coupled to at least one of the components. The example system further includes a sensor communication circuit that interprets a number of sensor data values in response to a sensed parameter group, a pattern recognition circuit that determines a recognized pattern value in response to a least a portion of the sensor data values, and a sensor learning circuit that updates the sensed parameter group in response to the recognized pattern value. The example sensor communication circuit further adjusts the interpreting of the sensor data values in response to the updated sensed parameter group.

[0993] Certain further aspects of an example system are described following, any one or more of which may be present in certain embodiments. An example system includes the sensed parameter group being a fused number of sensors, and where the recognized pattern value further includes a secondary value including a value determined in response to the fused number of sensors. An example system further includes the pattern recognition circuit and the sensor learning circuit iteratively performing the determining the recognized pattern value and the updating the sensed parameter group to improve a sensing performance value. An example system further includes the sensing performance value include a determination of one or more of the following: a signal-to-noise performance for detecting a value of interest in the industrial system; a network utilization of the sensors in the industrial system; an effective sensing resolution for a value of interest in the industrial system; a power consumption value for a sensing system in the industrial system, the sensing system including the sensors; a calculation efficiency for determining the secondary value; an accuracy and/or a precision of the secondary value; a redundancy capacity for determining the secondary value; and/or a lead time value for determining the secondary value. Example and non-

limiting calculation efficiency values include one or more determinations such as: processor operations to determine the secondary value; memory utilization for determining the secondary value; a number of sensor inputs from the number of sensors for determining the secondary value; and/or supporting data long-term storage for supporting the secondary value.

[0994] An example system includes one or more, or all, of the sensors as analog sensors and/or as remote sensors. An example system includes the secondary value being a value such as: a virtual sensor output value; a process prediction value; a process state value; a component prediction value; a component state value; and/or a model output value having the sensor data values from the fused number of sensors as an input. An example system includes the fused number of sensors being one or more of the combinations of sensors such as: a vibration sensor and a temperature sensor; a vibration sensor and a pressure sensor; a vibration sensor and an electric field sensor; a vibration sensor and a heat flux sensor; a vibration sensor and a galvanic sensor; and/or a vibration sensor and a magnetic sensor.

[0995] An example sensor learning circuit further updates the sensed parameter group by performing an operation such as: updating a sensor selection of the sensed parameter group; updating a sensor sampling rate of at least one sensor from the sensed parameter group; updating a sensor resolution of at least one sensor from the sensed parameter group; updating a storage value corresponding to at least one sensor from the sensed parameter group; updating a priority corresponding to at least one sensor from the sensed parameter group; and/or updating at least one of a sampling rate, sampling order, sampling phase, and/or a network path configuration corresponding to at least one sensor from the sensed parameter group. An example pattern recognition circuit further determines the recognized pattern value by performing an operation such as: determining a signal effectiveness of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to a value of interest; determining a sensitivity of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive confidence of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive delay time of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive accuracy of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive precision of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; and/or updating the recognized pattern value in response to external feedback. Example and non-limiting values of interest include: a virtual sensor output value; a process prediction value; a process state value; a component prediction value; a component state value; and/or a model output value having the sensor data values from the fused plurality of sensors as an input.

[0996] An example pattern recognition circuit further accesses cloud-based data including a second number of sensor data values, the second number of sensor data values corresponding to at least one offset industrial system. An example sensor learning circuit further accesses the cloud-

based data including a second updated sensor parameter group corresponding to the at least one offset industrial system.

[0997] An example procedure for data collection in an industrial environment includes an operation to provide a number of sensors to an industrial system including a number of components, each of the number of sensors operatively coupled to at least one of the number of components, an operation to interpret a number of sensor data values in response to a sensed parameter group, the sensed parameter group including a fused number of sensors from the number of sensors, an operation to determine a recognized pattern value including a secondary value determined in response to the number of sensor data values, an operation to update the sensed parameter group in response to the recognized pattern value, and an operation to adjust the interpreting the number of sensor data values in response to the updated sensed parameter group.

[0998] Certain further aspects of an example procedure are described following, any one or more of which may be included in certain embodiments. An example procedure includes an operation to iteratively perform the determining the recognized pattern value and the updating the sensed parameter group to improve a sensing performance value, where determining the sensing performance value includes an least one operation for determining a value, such as determining: a signal-to-noise performance for detecting a value of interest in the industrial system; a network utilization of the plurality of sensors in the industrial system; an effective sensing resolution for a value of interest in the industrial system; a power consumption value for a sensing system in the industrial system, the sensing system including the plurality of sensors; a calculation efficiency for determining the secondary value; an accuracy and/or a precision of the secondary value; a redundancy capacity for determining the secondary value; and/or a lead time value for determining the secondary value.

[0999] An example procedure includes an operation to update the sensed parameter group comprised by performing at least one operation such as: updating a sensor selection of the sensed parameter group; updating a sensor sampling rate of at least one sensor from the sensed parameter group; updating a sensor resolution of at least one sensor from the sensed parameter group; updating a storage value corresponding to at least one sensor from the sensed parameter group; updating a priority corresponding to at least one sensor from the sensed parameter group; and/or updating at least one of a sampling rate, sampling order, sampling phase, and a network path configuration corresponding to at least one sensor from the sensed parameter group. An example procedure includes determining the recognized pattern value by performing at least one operation such as: determining a signal effectiveness of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to a value of interest; determining a sensitivity of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive confidence of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive delay time of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive accuracy of at least one sensor of the sensed parameter group and the updated sensed parameter group

relative to the value of interest; determining a predictive precision of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; and/or updating the recognized pattern value in response to external feedback.

[1000] The term industrial system (and similar terms) as utilized herein should be understood broadly. Without limitation to any other aspect or description of the present disclosure, an industrial system includes any large scale process system, mechanical system, chemical system, assembly line, oil and gas system (including, without limitation, production, transportation, exploration, remote operations, offshore operations, and/or refining), mining system (including, without limitation, production, exploration, transportation, remote operations, and/or underground operations), rail system (yards, trains, shipments, etc.), construction, power generation, aerospace, agriculture, food processing, and/or energy generation. Certain components may not be considered industrial individually, but may be considered industrially in an aggregated system – for example a single fan, motor, and/or engine may be not an industrial system, but may be a part of a larger system and/or be accumulated with a number of other similar components to be considered an industrial system and/or a part of an industrial system. In certain example embodiments, a system may be considered an industrial system for some purposes but not for other purposes – for example a large data server farm may be considered an industrial system for certain sensing operations, such as temperature detection, vibration, or the like, but not an industrial system for other sensing operations such as gas composition. Additionally, in certain example embodiments, otherwise similar looking systems may be differentiated in determining whether such systems are industrial systems, and/or which type of industrial system. For example, one data server farm may not, at a given time, have process stream flow rates that are critical to operation, while another data server farm may have process stream flow rates that are critical to operation (e.g., a coolant flow stream), and accordingly one data farm server may be an industrial system for a data collection and/or sensing improvement process or system, while the other is not. Accordingly, the benefits of the present disclosure may be applied in a wide variety of systems, and any such systems may be considered an industrial system herein, while in certain embodiments a given system may not be considered an industrial system herein. One of skill in the art, having the benefit of the disclosure herein and knowledge about a contemplated system ordinarily available to that person, can readily determine which aspects of the present disclosure will benefit a particular system, how to combine processes and systems from the present disclosure to enhance operations of the contemplated system. Certain considerations for the person of skill in the art, in determining whether a contemplated system is an industrial system and/or whether aspects of the present disclosure can benefit or enhance the contemplated system include, without limitation: the accessibility of portions of the system to positioning sensing devices; the sensitivity of the system to capital costs (e.g., initial installation) and operating costs (e.g., optimization of processes, reduction of power usage); the transmission environment of the system (e.g., availability of broadband internet; satellite coverage; wireless cellular access; the electro-magnetic ("EM") environment of the system; the weather, temperature, and environmental conditions of the system; the availability of suitable locations to run wires, network lines, and the like; the presence and/or availability of

suitable locations for network infrastructure, router positioning, and/or wireless repeaters); the availability of trained personnel to interact with computing devices; the desired spatial, time, and/or frequency resolution of sensed parameters in the system; the degree to which a system or process is well understood or modeled; the turndown ratio in system operations (e.g., high load differential to low load; high flow differential to low flow; high temperature operation differential to low temperature operation); the turndown ratio in operating costs (e.g., effects of personnel costs based on time (day, season, etc.); effects of power consumption cost variance with time, throughput, etc.); the sensitivity of the system to failure, down-time, or the like; the remoteness of the contemplated system (e.g., transport costs, time delays, etc.); and/or qualitative scope of change in the system over the operating cycle (e.g., the system runs several distinct processes requiring a variable sensing environment with time; time cycle and nature of changes such as periodic, event driven, lead times generally available, etc.). While specific examples of industrial systems and considerations are described herein for purposes of illustration, any system benefitting from the disclosures herein, and any considerations understood to one of skill in the art having the benefit of the disclosures herein, are specifically contemplated within the scope of the present disclosure.

[1001] The term sensor (and similar terms) as utilized herein should be understood broadly. Without limitation to any other aspect or description of the present disclosure, sensor includes any device configured to provide a sensed value representative of a physical value (e.g., temperature, force, pressure) in a system, or representative of a conceptual value in a system at least having an ancillary relationship to a physical value (e.g., work, state of charge, frequency, phase, etc.).

[1002] Example and non-limiting sensors include vibration, acceleration, noise, pressure, force, position, location, velocity, displacement, temperature, heat flux, speed, rotational speed (e.g., a tachometer), motion, accelerometers, magnetic field, electrical field, galvanic, current, flow (gas, fluid, heat, particulates, particles, etc.), level, proximity, gas composition, fluid composition, toxicity, corrosiveness, acidity, pH, humidity, hygrometer measures, moisture, density (bulk or specific), ultrasound, imaging, analog, and/or digital sensors. The list of sensed values is a non-limiting example, and the benefits of the present disclosure in many applications can be realized independent of the sensor type, while in other applications the benefits of the present disclosure may be dependent upon the sensor type.

[1003] The sensor type and mechanism for detection may be any type of sensor understood in the art. Without limitation, an accelerometer may be any type and scaling, for example 500 mV per g (1 g = 9.8 m/s2), 100 mV, 1 V per g, 5 V per g, 10 V per g, 10 MV per g, as well as any frequency capability. It will be understood for accelerometers, and for all sensor types, that the scaling and range may be competing (e.g., in a fixed-bit or low bit A/D system), and/or selection of high resolution scaling with a large range may drive up sensor and/or computing costs, which may be acceptable in certain example embodiments, and may be prohibitive in other embodiments. Example and non-limiting accelerometers include piezo-electric devices, high resolution and sampling speed position detection devices (e.g., laser based devices), and/or detection of other parameters (strain, force, noise, etc.) that can be correlated to acceleration and/or vibration. Example and non-limiting proximity probes include electro-magnetic devices (e.g., Hall effect,

Variable Reluctance, etc.), a sleeve/oil film device, and/or determination of other parameters than can be correlated to proximity. An example vibration sensor includes a tri-axial probe, which may have high frequency response (e.g., scaling of 100 MV/g). Example and non-limiting temperature sensors include thermistors, thermocouples, and/or optical temperature determination.

[1004] A sensor may, additionally or alternatively, provide a processed value (e.g., a de-bounced, filtered, and/or compensated value) and/or a raw value, with processing downstream (e.g., in a data collector, controller, plant computer, and/or on a cloud-based data receiver). In certain example embodiments, a sensor provides a voltage, current, data file (e.g., for images), or other raw data output, and/or a sensor provides a value representative of the intended sensed measurement (e.g., a temperature sensor may communicate a voltage or a temperature value). Additionally or alternatively, a sensor may communicate wirelessly, through a wired connection, through an optical connection, or by any other mechanism. The described examples of sensor types and/or communication parameters are non-limiting examples for purposes of illustration.

[1005] Additionally or alternatively, in certain example embodiments, a sensor is a distributed physical device -- for example where two separate sensing elements coordinate to provide a sensed value (e.g., a position sensing element and a mass sensing element may coordinate to provide an acceleration value). In certain example embodiments, a single physical device may form two or more sensors, and/or parts of more than one sensor. For example, a position sensing element may form a position sensor and a velocity sensor, where the same physical hardware provides the sensed data for both determinations.

[1006] The term smart sensor, smart device (and similar terms) as utilized herein should be understood broadly. Without limitation to any other aspect or description of the present disclosure, a smart sensor includes any sensor and aspect thereof as described throughout the present disclosure. A smart sensor includes an increment of processing reflected in the sensed value communicated by the sensor, including at least basic sensor processing (e.g., de-bouncing, filtering, compensation, normalization, and/or output limiting), more complex compensations (e.g., correcting a temperature value based on known effects of current environmental conditions on the sensed temperature value, common mode or other noise removal, etc.), a sensing device that provides the sensed value as a network communication, and/or a sensing device that aggregates a number of sensed values for communication (e.g., multiple sensors on a device communicated out in a parseable or deconvolutable manner or as separate messages; multiple sensors providing a value to a single smart sensor, which relays sensed values on to a data collector, controller, plant computer, and/or cloud-based data receiver). The use of the term smart sensor is for purposes of illustration, and whether a sensor is a smart sensor can depend upon the context and the contemplated system, and can be a relative description compared to other sensors in the contemplated system. Thus, a given sensor having identical functionality may be a smart sensor for the purposes of one contemplated system, and just a sensor for the purposes of another contemplated system, and/or may be a smart sensor in a contemplated system during certain operating conditions, and just a sensor for the purposes of the same contemplated system during other operating conditions.

[1007] The terms sensor fusion, fused sensors, and similar terms, as utilized herein, should be understood broadly, except where context indicates otherwise, without limitation to any other aspect or description of the present disclosure. A sensor fusion includes a determination of second order data from sensor data, and further includes a determination of second order data from sensor data of multiple sensors, including involving multiplexing of streams of data, combinations of batches of data, and the like from the multiple sensors. Second order data include a determination about a system or operating condition beyond that which is sensed directly. For example, temperature, pressure, mixing rate, and other data may be analyzed to determine which parameters are result-effective on a desired outcome (e.g., a reaction rate). The sensor fusion may include sensor data from multiple sources, and/or longitudinal data (e.g., taken over a period of time, over the course of a process, and/or over an extent of components in a plant – for example tracking a number of assembled parts, a virtual slug of fluid passing through a pipeline, or the like). The sensor fusion may be performed in real-time (e.g., populating a number of sensor fusion determinations with sensor data as a process progresses), off-line (e.g., performed on a controller, plant computer, and/or cloud-based computing device), and/or as a post-processing operation (e.g., utilizing historical data, data from multiple plants or processes, etc.). In certain example embodiments, a sensor fusion includes a machine pattern recognition operation – for example where an outcome of a process is given to the machine and/or determined by the machine, and the machine pattern recognition operation determines result-effective parameters from the detected sensor value space to determine which operating conditions were likely to be the cause of the outcome and/or the off-nominal result of the outcome (e.g., process was less effective or more effective than nominal, failed, etc.). In certain example embodiments, the outcome may be a quantitative outcome (e.g., 20% more product was produced than a nominal run) or a qualitative outcome (e.g., product quality was unacceptable, component X of the contemplated system failed during the process, component X of the contemplated system required a maintenance or service event, etc.).

[1008] In certain example embodiments, a sensor fusion operation is iterative or recursive – for example an estimated set of result effective parameters is updated after the sensor fusion operation, and a subsequent sensor fusion operation is performed on the same data or another data set with an updated set of the result effective parameters. In certain example embodiments, subsequent sensor fusion operations include adjustments to the sensing scheme – for example higher resolution detections (e.g., in time, space, and/or frequency domains), larger data sets (and consequent commitment of computing and/or networking resources), changes in sensor capability and/or settings (e.g., changing an A/D scaling, range, resolution, etc.; changing to a more capable sensor and/or more capable data collector, etc.) are performed for subsequent sensor fusion operations. In certain example embodiments, the sensor fusion operation demonstrates improvements to the contemplated system (e.g., production quantity, quality, and/or purity, etc.) such that expenditure of additional resources to improve the sensing scheme are justified. In certain example embodiments, the sensor fusion operation provides for improvement in the sensing scheme without incremental cost – for example by narrowing the number of result effective parameters and thereby

freeing up system resources to provide greater resolution, sampling rates, etc., from hardware already present in the contemplated system. In certain example embodiments, iterative and/or recursive sensor fusion is performed on the same data set, a subsequent data set, and/or a historical data set. For example, high resolution data may already be present in the system, and a first sensor fusion operation is performed with low resolution data (e.g., sampled from the high resolution data set), such as to allow for completion of sensor fusion processing operations within a desired time frame, within a desired processor, memory, and/or network utilization, and/or to allow for checking a large number of variables as potential result effective parameters. In a further example, a greater number of samples from the high resolution data set may be utilized in a subsequent sensor fusion operation in response to confidence that improvements are present, narrowing of the potential result effective variables, and/or a determination that higher resolution data is required to determine the result effective parameters and/or effective values for such parameters.

[1009] The described operations and aspects for sensor fusion are non-limiting examples, and one of skill in the art, having the benefit of the disclosures herein and information ordinarily available about a contemplated system, can readily design a system to utilize and/or benefit from a sensor fusion operation. Certain considerations for a system to utilize and/or benefit from a sensor fusion operation include, without limitation: the number of components in the system; the cost of components in the system; the cost of maintenance and/or down-time for the system; the value of improvements in the system (production quantity, quality, yield, etc.); the presence, possibility, and/or consequences of undesirable system outcomes (e.g., side products, thermal and/or luminary events, environmental benefits or consequences, hazards present in the system); the expense of providing a multiplicity of sensors for the system; the complexity between system inputs and system outputs; the availability and cost of computing resources (e.g., processing, memory, and/or communication throughput); the size/scale of the contemplated system and/or the ability of such a system to generate statistically significant data; whether offset systems exist, including whether data from offset systems is available and whether combining data from offset systems will generate a statistically improved data set relative to the system considered alone; and/or the cost of upgrading, improving, or changing a sensing scheme for the contemplated system. The described considerations for a contemplated system that may benefit from or utilize a sensor fusion operation are non-limiting illustrations.

[1010] Certain systems, processes, operations, and/or components are described in the present disclosure as "offset systems" or the like. An offset system is a system distinct from a contemplated system, but having relevance to the contemplated system. For example, a contemplated refinery may have an "offset refinery," which may be a refinery operated by a competitor, by a same entity operating the contemplated refinery, and/or a historically operated refinery that no longer exists. The offset refinery bears some relevant relationship to the contemplated refinery, such as utilizing similar reactions, process flows, production volumes, feed stock, effluent materials, or the like. A system which is an offset system for one purpose may not be an offset system for another purpose. For example, a manufacturing process utilizing conveyor belts and similar motors may be an offset process for a contemplated manufacturing process for the purpose of tracking product movement,

understanding motor operations and failure modes, or the like, but may not be an offset process for product quality if the products being produced have distinct quality outcome parameters. Any industrial system contemplated herein may have an offset system for certain purposes. One of skill in the art, having the benefit of the present disclosure and information ordinarily available for a contemplated system, can readily determine what is disclosed by an offset system or offset aspect of a system.

[1011] Any one or more of the terms computer, computing device, processor, circuit, and/or server include a computer of any type, capable to access instructions stored in communication thereto such as upon a non-transient computer readable medium, whereupon the computer performs operations of systems or methods described herein upon executing the instructions. In certain example embodiments, such instructions themselves comprise a computer, computing device, processor, circuit, and/or server. Additionally or alternatively, a computer, computing device, processor, circuit, and/or server may be a separate hardware device, one or more computing resources distributed across hardware devices, and/or may include such aspects as logical circuits, embedded circuits, sensors, actuators, input and/or output devices, network and/or communication resources, memory resources of any type, processing resources of any type, and/or hardware devices configured to be responsive to determined conditions to functionally execute one or more operations of systems and methods herein.

[1012] Certain operations described herein include interpreting, receiving, and/or determining one or more values, parameters, inputs, data, or other information. Operations including interpreting, receiving, and/or determining any value parameter, input, data, and/or other information include, without limitation: receiving data via a user input; receiving data over a network of any type; reading a data value from a memory location in communication with the receiving device; utilizing a default value as a received data value; estimating, calculating, or deriving a data value based on other information available to the receiving device; and/or updating any of these in response to a later received data value. In certain example embodiments, a data value may be received by a first operation, and later updated by a second operation, as part of the receiving a data value. For example, when communications are down, intermittent, or interrupted, a first operation to interpret, receive, and/or determine a data value may be performed, and when communications are restored an updated operation to interpret, receive, and/or determine the data value may be performed.

[1013] Certain logical groupings of operations herein, for example methods or procedures of the current disclosure, are provided to illustrate aspects of the present disclosure. Operations described herein are schematically described and/or depicted, and operations may be combined, divided, re-ordered, added, or removed in a manner consistent with the disclosure herein. It is understood that the context of an operational description may require an ordering for one or more operations, and/or an order for one or more operations may be explicitly disclosed, but the order of operations should be understood broadly, where any equivalent grouping of operations to provide an equivalent outcome of operations is specifically contemplated herein. For example, if a value is used in one operational step, the determining of the value may be required before that operational step in certain contexts (e.g., where the time delay of data for an operation to achieve a certain effect is important),

but may not be required before that operation step in other contexts (e.g., where usage of the value from a previous execution cycle of the operations would be sufficient for those purposes). Accordingly, in certain embodiments an order of operations and grouping of operations as described is explicitly contemplated herein, and in certain embodiments re-ordering, subdivision, and/or different grouping of operations is explicitly contemplated herein.

[1014] Referencing Figure 104, an example system 10902 for data collection in an industrial environment includes an industrial system 10904 having a number of components 10906, and a number of sensors 10908, wherein each of the sensors 10908 is operatively coupled to at least one of the components 10906. The selection, distribution, type, and communicative setup of sensors depends upon the application of the system 10902 and/or the context.

[1015] The example system 10902 further includes a sensor communication circuit 10920 (reference Figure 105) that interprets a number of sensor data values 10948 in response to a sensed parameter group 10928. The sensed parameter group 10928 includes a description of which sensors 10908 are sampled at which times, including at least the selected sampling frequency, a process stage wherein a particular sensor may be providing a value of interest, and the like. An example system includes the sensed parameter group 10928 being a fused number of sensors 10926, for example a set of sensors believed to encompass detection of operating conditions of the system that affect a desired output, such as production output, quality, efficiency, profitability, purity, maintenance or service predictions of components in the system, failure mode predictions, and the like. In a further embodiment, the recognized pattern value 10930 further includes a secondary value 10932 including a value determined in response to the fused number of sensors 10926.

[1016] In certain example embodiments, sensor data values 10948 are provided to a data collector 10910, which may be in communication with multiple sensors 10908 and/or with a controller 10914. In certain example embodiments, a plant computer 10912 is additionally or alternatively present. In the example system, the controller 10914 is structured to functionally execute operations of the sensor communication circuit 10920, pattern recognition circuit 10922, and/or the sensor learning circuit 10924, and is depicted as a separate device for clarity of description. Aspects of the controller 10914 may be present on the sensors 10908, the data controller 10910, the plant computer 10912, and/or on a cloud computing device 10916. In certain example embodiments, all aspects of the controller 10914 may be present in another device depicted on the system 10902. The plant computer 10912 represents local computing resources, for example processing, memory, and/or network resources, that may be present and/or in communication with the industrial system 10904. In certain example embodiments, the cloud computing device 10916 represents computing resources externally available to the industrial system 10904, for example over a private network, intra-net, through cellular communications, satellite communications, and/or over the internet. In certain example embodiments, the data controller 10910 may be a computing device, a smart sensor, a MUX box, or other data collection device capable to receive data from multiple sensors and to pass-through the data and/or store data for later transmission. An example data controller 10910 has no storage and/or limited storage, and selectively passes sensor data therethrough, with a subset of the sensor data being communicated at a given time due to

bandwidth considerations of the data controller 10910, a related network, and/or imposed by environmental constraints. In certain example embodiments, one or more sensors and/or computing devices in the system 10902 are portable devices – for example a plant operator walking through the industrial system may have a smart phone, which the system 10902 may selectively utilize as a data controller 10910, sensor 10908 – for example to enhance communication throughput, sensor resolution, and/or as a primary method for communicating sensor data values 10948 to the controller 10914.

[1017] The example system 10902 further includes a pattern recognition circuit 10922 that determines a recognized pattern value 10930 in response to a least a portion of the sensor data values 10948.

[1018] The example system 10902 further includes a sensor learning circuit 10924 that updates the sensed parameter group 10928 in response to the recognized pattern value 10930. The example sensor communication circuit 10920 further adjusts the interpreting of the sensor data values 10948 in response to the updated sensed parameter group 10928.

[1019] An example system 10902 further includes the pattern recognition circuit 10922 and the sensor learning circuit 10924 iteratively performing the determining the recognized pattern value 10930 and the updating the sensed parameter group 10928 to improve a sensing performance value 10934. For example, the pattern recognition circuit 10922 may add sensors, remove sensors, and/or change sensor setting to modify the sensed parameter group 10928 based upon sensors which appear to be effective or ineffective predictors of the recognized pattern value 10930, and the sensor learning circuit 10924 may instruct a continued change (e.g., while improvement is still occurring), an increased or decreased rate of change (e.g., to converge more quickly on an improved sensed parameter group 10928), and/or instruct a randomized change to the sensed parameter group 10928 (e.g., to ensure that all potentially result effective sensors are being checked, and/or to avoid converging into a local optimal value).

[1020] Example and non-limiting options for the sensing performance value 10934 include: a signal-to-noise performance for detecting a value of interest in the industrial system (e.g., a determination that the prediction signal for the value is high relative to noise factors for one or more sensors of the sensed parameter group 10928, and/or for the sensed parameter group 10928 as a whole); a network utilization of the sensors in the industrial system (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928 relatively high where it is as effective or almost as effective as another sensed parameter group 10928, but results in lower network utilization); an effective sensing resolution for a value of interest in the industrial system (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928 relatively high where it provides a responsive prediction of the output value to smaller changes in input values); a power consumption value for a sensing system in the industrial system, the sensing system including the sensors (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928 relatively high where it is as effective or almost as effective as another sensed parameter group 10928, but results in lower power consumption) ; a calculation efficiency for determining the secondary value (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928

relatively high where it is as effective or almost as effective as another sensed parameter group 10928 in determining the secondary value 10932, but results in fewer processor cycles, lower network utilization, and/or lower memory utilization including stored memory requirements as well as intermediate memory utilization such as buffers); an accuracy and/or a precision of the secondary value (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928 relatively high where it provides a highly accurate and/or highly precise determination of the secondary value 10932); a redundancy capacity for determining the secondary value (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928 relatively high where it provides similar capability and/or resource utilization, but provides for additional sensing redundancy, such as being more robust to gaps in data from one or more of the sensors in the sensed parameter group 10928); and/or a lead time value for determining the secondary value 10932 (e.g., the sensor learning circuit 10924 may score a sensed parameter group 10928 relatively high where it provides an improved or sufficient lead time in the secondary value 10932 determination – for example to assist in avoiding over-temperature operation, spoiling an entire production run, determining whether a component has sufficient service life to complete a production run, etc.) Example and non-limiting calculation efficiency values include one or more determinations such as: processor operations to determine the secondary value 10932; memory utilization for determining the secondary value 10932; a number of sensor inputs from the number of sensors for determining the secondary value 10932; and/or supporting memory, such as long-term storage or buffers for supporting the secondary value 10932.

[1021] Example systems include one or more, or all, of the sensors 10908 as analog sensors and/or as remote sensors. An example system includes the secondary value 10932 being a value such as: a virtual sensor output value; a process prediction value (e.g., a success value for a production run, an overtemperature value, an overpressure value, a product quality value, etc.); a process state value (e.g., a stage of the process, a temperature at a time and location in the process); a component prediction value (e.g., a component failure prediction, a component maintenance or service prediction, a component response to an operating change prediction); a component state value (a remaining service life or maintenance interval for a component); and/or a model output value having the sensor data values 10948 from the fused number of sensors 10926 as an input. An example system includes the fused number of sensors 10926 being one or more of the combinations of sensors such as: a vibration sensor and a temperature sensor; a vibration sensor and a pressure sensor; a vibration sensor and an electric field sensor; a vibration sensor and a heat flux sensor; a vibration sensor and a galvanic sensor; and/or a vibration sensor and a magnetic sensor.

[1022] An example sensor learning circuit 10924 further updates the sensed parameter group 10928 by performing an operation such as: updating a sensor selection of the sensed parameter group 10928 (e.g., which sensors are sampled); updating a sensor sampling rate of at least one sensor from the sensed parameter group (e.g., how fast the sensors provide information, and/or how fast information is passed through the network); updating a sensor resolution of at least one sensor from the sensed parameter group (e.g., changing or requesting a change in a sensor resolution, utilizing additional sensors to provide greater effective resolution); updating a storage

value corresponding to at least one sensor from the sensed parameter group (e.g., storing data from the sensor at a higher or lower resolution, and/or over a longer or shorter time period); updating a priority corresponding to at least one sensor from the sensed parameter group (e.g., moving a sensor up to a higher priority – for example, if environmental conditions prevent data receipt from all planned sensors, and/or reducing a time lag between creation of the sensed data and receipt at the sensor learning circuit 10924); and/or updating at least one of a sampling rate, sampling order, sampling phase, and/or a network path configuration corresponding to at least one sensor from the sensed parameter group.

[1023] An example pattern recognition circuit 10922 further determines the recognized pattern value 10930 by performing an operation such as: determining a signal effectiveness of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to a value of interest 10950 (e.g., determining that a sensor value is a good predictor of the value of interest 10950); determining a sensitivity of at least one sensor of the sensed parameter group 10928 and the updated sensed parameter group 10928 relative to the value of interest 10950 (e.g., determining the relative sensitivity of the determined value of interest to small changes in operating conditions based on the selected sensed parameter group 10928); determining a predictive confidence of at least one sensor of the sensed parameter group 10928 and the updated sensed parameter group 10928 relative to the value of interest 10950; determining a predictive delay time of at least one sensor of the sensed parameter group 10928 and the updated sensed parameter group 10928 relative to the value of interest 10950; determining a predictive accuracy of at least one sensor of the sensed parameter group 10928 and the updated sensed parameter group 10928 relative to the value of interest 10950; determining a classification precision of at least one sensor of the sensed parameter group 10928 (e.g., determining the accuracy of classification of a pattern by a machine classifier based on use of the at least one sensor); determining a predictive precision of at least one sensor of the sensed parameter group 10928 and the updated sensed parameter group 10928 relative to the value of interest 10950; and/or updating the recognized pattern value 10930 in response to external feedback, which may be received as external data 10952 (e.g., where an outcome is known, such as a maintenance event, product quality determination, production outcome determination, etc., the detection of the recognized pattern value 10930 is thereby improved according to the conditions of the system before the known outcome occurred). Example and non-limiting values of interest 10950 include: a virtual sensor output value; a process prediction value; a process state value; a component prediction value; a component state value; and/or a model output value having the sensor data values from the fused plurality of sensors as an input.

[1024] An example pattern recognition circuit 10922 further accesses cloud-based data 10954 including a second number of sensor data values, the second number of sensor data values corresponding to at least one offset industrial system. An example sensor learning circuit 10924 further accesses the cloud-based data 10954 including a second updated sensor parameter group corresponding to the at least one offset industrial system. Accordingly, the pattern recognition circuit 10922 can improve pattern recognition in the system based on increased statistical data available from an offset system. Additionally, or alternatively, the sensor learning circuit 10924

can improve more rapidly and with greater confidence based upon the data from the offset system – including determining which sensors in the offset system found to be effective in predicting system outcomes.

[1025] An example system includes an industrial system including an oil refinery. An example oil refinery includes one or more compressors for transferring fluids throughout the plant, and/or for pressurizing fluid streams (e.g., for reflux in a distillation column). Additionally, or alternatively, the example oil refinery includes vacuum distillation, for example, to fractionate hydrocarbons. The example oil refinery additionally includes various pipelines in the system for transferring fluids, bringing in feedstock, final product delivery, and the like. An example system includes a number of sensors configured to determine each aspect of a distillation column – for example temperatures of various fluid streams, temperatures, and compositions of individual contact trays in the column, measurements of the feed and reflux, as well as of the effluent or separated products. The design of a distillation column is complex, and optimal design can depend upon the sizing of boilers, compressors, the contact conditions within the column, as well as the composition of feedstock, all of which can vary significantly. Additionally, the optimal position for effective sensing of conditions in a pipeline can vary with fluid flow rates, environmental conditions (e.g., causing variation in heat transfer rates), the feedstock utilized, and other factors. Additionally, wear or loss of capability in a boiler, compressor, or other operating equipment can change the system response and capabilities, rendering a single point optimization – including where sensors should be positioned and how they should sample data – to be non-optimal as the system ages.

[1026] Provision of multiple sensors throughout the system can be costly, not necessarily because the sensors are expensive, but because the sensors provide data which may be prohibitive to transmit, store, and utilize. Cost may involve costs of transmitting over networks, as well as costs of operations, such as numbers of input/output operations (and time required to undertake such operations). The example system includes providing a large number of sensors throughout the system, and determining which of the sensors are effective for control and optimization of the distillation process. Additionally, as the feedstock and/or environmental conditions change, the optimal sensor package for both optimization and control may change. The example system utilizes a pattern recognition circuit to determine which sensors, including sensor fusion operations (including selection of groups, selection of multiplexing and combination, and the like), are effective in controlling the desired parameters of the distillation, and in determining the optimal values for temperatures, flow rates, entry trays for feed and reflux, and/or reflux rates. Additionally, the sensor learning circuit is capable, over time and/or utilizing offset oil refineries, to rapidly converge on various sensor packages that are appropriate for a multiplicity of operating conditions. If an unexpected operating condition occurs – for example an off-nominal operation of a compressor, the sensor learning circuit is capable of migrating the system to the correct sensing and operating conditions for the unexpected operating condition. The ability to flexibly utilize a multiplicity of sensors allows for the system to be flexible in response to changing conditions without providing for excessive capability in transmission and storage of sensor data. Accordingly, operations of the distillation column are improved and can be optimized for a large number of

operating conditions. Additionally, alerts for the distillation column, based upon recognition of patterns indicating off-nominal operation, can be readily prepared to adjust or shut down the process before significant product quality loss and/or hazardous conditions develop. Example sensor fusion operations for a refinery include vibration information combined with temperatures, pressures, and/or composition (e.g., to determine compressor performance); temperature and pressure, temperature and composition, and/or composition and pressure (e.g., to determine feedstock variance, contact tray performance, and/or a component failure).

[1027] An example refinery system includes storage tanks and/or boiler feed water. Example system determinations include a sensor fusion to determine a storage tank failure and/or off-nominal operation, such as through a temperature and pressure fusion, and/or a vibration determination with a non-vibration determination (e.g., detecting leaks, air in the system, and/or a feed pump issue). Certain further example system determinations include a sensor fusion to determine a boiler feed water failure, such as through a sensor fusion including flow rate, pressure, temperature, and/or vibration. Any one or more of these parameters can be utilized to determine a system leak, failure, wear of a feed pump, scaling, and/or to reduce pumping losses while maintaining system flow rates. Similarly, an example industrial system includes a power generation system having a condensate and/or make-up water system, where a sensor fusion provides for a sensed parameter group and prediction of failures, maintenance, and the like.

[1028] An example industrial system includes an irrigation system for a field or a system of fields. Irrigations systems are subject to significant variability in the system (e.g., inlet pressures and/or water levels, component wear and maintenance) as well as environmental variability (e.g., types and distribution of crops planted, weather, soil moisture, humidity, seasonal variability in the sun, cloud coverage, and/or wind variance). Additionally, irrigation systems tend to be remotely located where high bandwidth network access, maintenance facilities, and/or even personnel for oversight are not readily available. An example system includes a multiplicity of sensors capable of detecting conditions for the irrigation system, without requiring that all of the sensors transmit or store data on a continuous basis. The pattern recognition circuit can readily determine the most important set of sensors to effectively predict patterns and those system conditions requiring a response (e.g., irrigation cycles, positioning, and the like). The sensor learning circuit provides for responsive migration of the sensed parameter group to variability, which may occur on slower (e.g., seasonal, climate change, etc.) or faster cycles (e.g., equipment failure, weather conditions, step change events such as planting or harvesting). Additionally, alerts for remote facilities can be readily prepared with confidence that the correct sensor package is in place for determining an off-nominal condition (e.g., imminent failure or maintenance requirement for a pump).

[1029] An example industrial system includes a chemical or pharmaceutical plant. Chemical plants require specific operating conditions, flow rates, temperatures, and the like to maintain proper temperatures, concentrations, mixing, and the like throughout the system. In many systems, there are numerous process steps, and an off-nominal or uncoordinated operation in one part of the process can result in reduced yields, a failed process, and/or a significant reduction in production capacity as coordinated processes must respond (or as coordinated processes fail to respond).

Accordingly, a very large number of systems are required to minimally define the system, and in certain embodiments a prohibitive number of sensors are required, from a data transmission and storage viewpoint, to keep sensing capability for a broad range of operating conditions. Additionally, the complexity of the system results in difficulty optimizing and coordinating system operations even where sufficient sensors are present. In certain example embodiments, the pattern recognition circuit can determine the sensing parameter groups that provide high resolution understanding of the system, without requiring that all of the sensors store and transmit data continuously. Further, the utilization of a sensor fusion provides for the opportunity to abstract desired outputs, for example "maximize yield" or "minimize an undesirable side reaction" without requiring a full understanding from the operator of which sensors and system conditions are most effective to achieve the abstracted desired output. Example components in a chemical or pharmaceutical plan amenable to control and predictions based on a sensor fusion operation include an agitator, a pressure reactor, a catalytic reactor, and/or a thermic heating system. Example sensor fusion operations to determine sensed parameter groups and tune the pattern recognition circuit include, without limitation, a vibration sensor combined with another sensor type, a composition sensor combined with another sensor type, a flow rate determination combined with another sensor type, and/or a temperature sensor combined with another sensor type. The sensor fusion best suited for a particular application can be converged upon by the sensor learning circuit, but also depends upon the type of component that is subject to predictions, as well as the type of desired outputs pursued by the operator. For example, agitators are amenable to vibration sensing, as well as uniformity of composition detection (e.g., high resolution temperature), expected reaction rates in a properly mixed system, and the like. Catalytic reactors are amenable to temperature sensing (based on the reaction thermodynamics), composition detection (e.g., for expected reactants, as well as direct detection of catalytic material), flow rates (e.g., gross mechanical failure, reduced volume of beads, etc.), and/or pressure detection (e.g., indicative of or coupled with flow rate changes).

[1030] An example industrial system includes a food processing system. Example food processing systems include pressurization vessels, stirrers, mixers, and/or thermic heating systems. Control of the process is critical to maintain food safety, product quality, and product consistency. However, most input parameters to the food processing system are subject to high variability – for example basic food products are inherently variable as natural products, with differing water content, protein content, and aesthetic variation. Additionally, labor cost management, power cost management, and variability in supply water, etc., provide for a complex process where determination of the process control variables, sensed parameters to determine these, and optimization of sensing in response to process variation are a difficult problem to resolve. Food processing systems are often cost conscious, and capital costs (e.g., for a robust network and computing system for optimization) are not readily incurred. Further, a food processing system may manufacture a wide variety of products on similar or the same production facilities – for example, to support an entire product line and/or due to seasonal variations. Accordingly, a sensor setup for one process may not support another process well. An example system includes the pattern recognition circuit determining the

sensing parameter groups that provide a strong signal response in target outcomes even in light of high variability in system conditions. The pattern recognition circuit can provide for numerous sensed group parameter options available for different process conditions without requiring extensive computing or data storage resources. Additionally, the sensor learning circuit provides for rapid response of the sensing system to changes in the process conditions, including updating the sensed group parameter options to pursue abstracted target outputs without the operator having to understand which sensed parameters best support the output goals. The sensor fusion best suited for a particular application can be converged upon by the sensor learning circuit, but also depends upon the type of component that is subject to predictions, as well as the type of desired outputs pursued by the operator. For example, control of and predictions for pressurization vessels, stirrers, mixers, and/or thermic heating systems are amenable to a sensor fusion with a temperature determination combined with a non-temperature determination, a vibration determination combined with a non-vibration determination, and/or a heat map combined with a rate of change in the heat map and/or a non-heat map determination. An example system includes a sensor fusion with a vibration determination and a non-vibration determination, wherein predictive information for a mixer and/or a stirrer is provided. An example system includes a sensor fusion with a pressure determination, a temperature determination, and/or a non-pressure determination, wherein predictive information for a pressurization vessel is provided.

[1031] Referencing Figure 106, an example procedure 10936 for data collection in an industrial environment includes an operation 10938 to provide a number of sensors to an industrial system including a number of components, each of the number of sensors operatively coupled to at least one of the number of components. The procedure 10936 further includes an operation 10940 to interpret a number of sensor data values in response to a sensed parameter group, the sensed parameter group including a fused number of sensors from the number of sensors, an operation 10942 to determine a recognized pattern value including a secondary value determined in response to the number of sensor data values, an operation 10944 to update the sensed parameter group in response to the recognized pattern value, and an operation 10946 to adjust the interpreting the number of sensor data values in response to the updated sensed parameter group.

[1032] An example procedure 10936 includes an operation to iteratively perform the determining the recognized pattern value and the updating the sensed parameter group to improve a sensing performance value (e.g., by repeating operations 10940 to 10944 periodically, at selected intervals, and/or in response to a system change). An example procedure 10936 includes determining the sensing performance value by determining: a signal-to-noise performance for detecting a value of interest in the industrial system; a network utilization of the plurality of sensors in the industrial system; an effective sensing resolution for a value of interest in the industrial system; a power consumption value for a sensing system in the industrial system, the sensing system including the plurality of sensors; a calculation efficiency for determining the secondary value; an accuracy and/or a precision of the secondary value; a redundancy capacity for determining the secondary value; and/or a lead time value for determining the secondary value.

[1033] An example procedure 10936 includes the operation 10944 to update the sensed parameter group by performing at least one operation such as: updating a sensor selection of the sensed parameter group; updating a sensor sampling rate of at least one sensor from the sensed parameter group; updating a sensor resolution of at least one sensor from the sensed parameter group; updating a storage value corresponding to at least one sensor from the sensed parameter group; updating a priority corresponding to at least one sensor from the sensed parameter group; and/or updating at least one of a sampling rate, sampling order, sampling phase, and a network path configuration corresponding to at least one sensor from the sensed parameter group. An example procedure 10936 includes the operation 10942 to determine the recognized pattern value by performing at least one operation such as: determining a signal effectiveness of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to a value of interest; determining a sensitivity of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive confidence of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive delay time of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive accuracy of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive precision of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; and/or updating the recognized pattern value in response to external feedback.

[1034] Referencing Figure 107, an example system 11000 for data collection in an industrial environment includes an industrial system 11002 having a number of components 11004, and a number of sensors 11006 each operatively coupled to at least one of the number of components 11004. The selection, distribution, type, and communicative setup of sensors depends upon the application of the system 11000 and/or the context.

[1035] The example system 11000 further includes a sensor communication circuit 11018 (reference Figure 108) that interprets a number of sensor data values 11034 in response to a sensed parameter group 11026. The sensed parameter group 11026 includes a description of which sensors 11006 are sampled at which times, including at least the selected sampling frequency, a process stage wherein a particular sensor may be providing a value of interest, and the like. An example system includes the sensed parameter group 11026 being a number of sensors provided for a sensor fusion operation. In certain example embodiments, the sensed parameter group 11026 includes a set of sensors that encompass detection of operating conditions of the system that predict outcomes, off-nominal operations, maintenance intervals, maintenance health states, and/or future state values for any of these, for a process, a component, a sensor, and/or any aspect of interest for the system 11000.

[1036] In certain example embodiments, sensor data values 11034 are provided to a data collector 11008, which may be in communication with multiple sensors 11006 and/or with a controller 11012. In certain example embodiments, a plant computer 11010 is additionally or alternatively

present. In the example system, the controller 11012 is structured to functionally execute operations of the sensor communication circuit 11018, pattern recognition circuit 11020, and/or the system characterization circuit 11022, and is depicted as a separate device for clarity of description. Aspects of the controller 11012 may be present on the sensors 11006, the data collector 11008, the plant computer 11010, and/or on a cloud computing device 11014. In certain example embodiments, all aspects of the controller 11012 may be present in another device depicted on the system 11000. The plant computer 11010 represents local computing resources, for example processing, memory, and/or network resources, that may be present and/or in communication with the industrial system 11000. In certain example embodiments, the cloud computing device 11014 represents computing resources externally available to the industrial system 11000, for example over a private network, intra-net, through cellular communications, satellite communications, and/or over the internet. In certain example embodiments, the data collector 11008 may be a computing device, a smart sensor, a MUX box, or other data collection device capable to receive data from multiple sensors and to pass-through the data and/or store data for later transmission. An example data collector 11008 has no storage and/or limited storage, and selectively passes sensor data therethrough, with a subset of the sensor data being communicated at a given time due to bandwidth considerations of the data collector 11008, a related network, and/or imposed by environmental constraints. In certain example embodiments, one or more sensors and/or computing devices in the system 11000 are portable devices – for example a plant operator walking through the industrial system may have a smart phone, which the system 11000 may selectively utilize as a data collector 11008, sensor 11006 – for example to enhance communication throughput, sensor resolution, and/or as a primary method for communicating sensor data values 11034 to the controller 11012.

[1037] The example system 11000 further includes a pattern recognition circuit 11020 that determines a recognized pattern value 11028 in response to a least a portion of the sensor data values 11034, and a system characterization circuit 11022 that provides a system characterization value 11030 for the industrial system in response to the recognized pattern value 11028. The system characterization value 11030 includes any value determined from the pattern recognition operations of the pattern recognition circuit 11020, including determining that a system condition of interest is present, a component condition of interest is present, an abstracted condition of the system or a component is present (e.g., a product quality value; an operation cost value; a component health, wear, or maintenance value; a component capacity value; and/or a sensor saturation value) and/or is predicted to occur within a time frame (e.g., calendar time, operational time, and/or a process stage) of interest. Pattern recognition operations include determining that operations compatible with a previously known pattern, operations similar to a previously known pattern and/or extrapolated from previously known pattern information (e.g., a previously known pattern includes a temperature response for a first component, and a known or estimated relationship between components allows for a determination that a temperature for a second component will exceed a threshold based upon the pattern recognition for the first component combined with the known or estimated relationship).

[1038] Non-limiting descriptions of a number of examples of a system characterization value 11030 are described following. An example system characterization value 11030 includes a predicted outcome for a process associated with the industrial system – for example a product quality description, a product quantity description, a product variability description (e.g., the expected variability of a product parameter predicted according to the operating conditions of the system), a product yield description, a net present value (NPV) for a process, a process completion time, a process chance of completion success, and/or a product purity result. The predicted outcome may be a batch prediction (e.g., a single run, or an integer number of runs, of the process, and the associated predicted outcome), a time based prediction (e.g., the projected outcome of the process over the next day, the next three weeks, until a scheduled shutdown, etc.), a production defined prediction (e.g., the projected outcome over the next 1,000 units, over the next 47 orders, etc.), and/or a rate of change based outcome (e.g., projected for 3 component failures per month, an emissions output per year, etc.). An example system characterization value 11030 includes a predicted future state for a process associated with the industrial system – for example an operating temperature at a given future time, an energy consumption value, a volume in a tank, an emitted noise value at a school adjacent to the industrial system, and/or a rotational speed of a pump. The predicted future state may be time based (e.g., at 4 PM on Thursday), based on a state of the process (e.g., during the third stage, during system shutdown, etc.), and/or based on a future state of particular interest (e.g., peak energy consumption, highest temperature value, maximum noise value, time or process stage when a maximum number of personnel will be within 50 feet of a sensitive area, time or process stage when an aspect of the system redundancy is at a lowest point – e.g., for determining high risk points in a process, etc.). An example system characterization value 11030 includes a predicted off-nominal operation for the process associated with the industrial system – for example when a component capacity of the system will exceed nominal parameters (although, possibly, not experience a failure), when any parameter in the system will be three standard deviations away from normal operations, when a capacity of a component will be under-utilized, etc. An example system characterization value 11030 includes a prediction value for one of the number of components – for example an operating condition at a point in time and/or process stage. An example system characterization value 11030 includes a future state value for one of the number of components. The predicted future state of a component may be time based, based on a state of the process, and/or based on a future state of particular interest (e.g., a highest or lowest value predicted for the component). An example system characterization value 11030 includes an anticipated maintenance health state information for one of the number of components, including at a particular time, a process stage, a lowest value predicted until a next maintenance event, etc. An example system characterization value 11030 includes a predicted maintenance interval for at least one of the number of components (e.g., based on current usage, anticipated usage, planned process operations, etc.). An example system characterization value 11030 includes a predicted off-nominal operation for one of the number of components – for example at a selected time, a process stage, and/or a future state of particular interest. An example system characterization value 11030 includes a predicted fault operation for one of the plurality of

components – for example at a selected time, a process stage, any fault occurrence predicted based on current usage, anticipated usage, planned process operations, and/or a future state of particular interest. An example system characterization value 11030 includes a predicted exceedance value for one of the number of components, where the exceedance value includes exceedance of a design specification, and/or exceedance of a selected threshold. An example system characterization value 11030 includes a predicted saturation value for one of the plurality of sensors for example at a selected time, a process stage, any saturation occurrence predicted based on current usage, anticipated usage, planned process operations, and/or a future state of particular interest.

[1039] Any values for the prediction value 11030 may be raw values (e.g., a temperature value), derivative values (e.g., a rate of change of a temperature value), accumulated values (e.g., a time spent above one or more temperature thresholds) including weighted accumulated values, and/or integrated values (e.g., an area over a temperature-time curve at a temperature value or temperature trajectory of interest). The provided examples list temperature, but any prediction value 11030 may be utilized, including at least vibration, system throughput, pressure, etc. In certain example embodiments, combinations of one or more prediction values 11030 may be utilized.

[1040] It will be appreciated in light of the disclosure that combining prediction values 11030 can create particularly powerful combinations for system analysis, control, and risk management, which are specifically contemplated herein. For example, a first prediction value may indicate a time or process stage for a maximum flow rate through the system, and a second prediction value may determine the predicted state of one or more components of the system that is present at that particular time or process stage. In another example, a first prediction value indicates a lowest margin of the system in terms of capacity to deliver (e.g., by determining a point in the process wherein at least one component has a lowest operating margin, and/or where a group of components have a statistically lower operating margin due to the risk induced by a number of simultaneous low operating margins), and a second prediction value testing a system risk (e.g., loss of inlet water, loss of power, increase in temperature, change in environmental conditions that reduce or increase heat transfer, or that preclude the emission of certain effluents), and the combined risk of separate events can be assessed on the total system risk. Additionally, the prediction values may be operated with a sensitivity check (e.g., varying system conditions within margins to determine if some failure may occur), wherein the use of the prediction value allows for the sensitivity check to be performed with higher resolution at high risk points in the process.

[1041] An example system 11000 further includes a system collaboration circuit 11024 that interprets external data 11036, and where the pattern recognition circuit 11020 further determines the recognized pattern value 11028 further in response to the external data 11036. External data 11036 includes, without limitation, data provided from outside the system 11000 and/or outside the controller 11012. Non-limiting example external data 11036 include entries from an operator (e.g., indicating a failure, a fault, and/or a service event). An example pattern recognition circuit 11020 further iteratively improves pattern recognition operations in response to the external data 11036 (e.g., where an outcome is known, such as a maintenance event, product quality determination, production outcome determination, etc., the detection of the recognized pattern

value 11028 is thereby improved according to the conditions of the system before the known outcome occurred). Example and non-limiting external data 11036 include data such as: an indicated process success value; an indicated process failure value; an indicated component maintenance event; an indicated component failure event; an indicated process outcome value; an indicated component wear value; an indicated process operational exceedance value; an indicated component operational exceedance value; an indicated fault value; and/or an indicated sensor saturation value.

[1042] An example system 11000 further includes a system collaboration circuit 11024 that interprets cloud-based data 11032 including a second number of sensor data values, the second number of sensor data values corresponding to at least one offset industrial system, and where the pattern recognition circuit 11020 further determines the recognized pattern value 11028 further in response to the cloud-based data 11032. An example pattern recognition circuit 11020 further iteratively improves pattern recognition operations in response to the cloud-based data 11032. An example sensed parameter group 11026 includes a triaxial vibration sensor, a vibration sensor and a second sensor that is not a vibration sensor, the second sensor being a digital sensor, and/or a number of analog sensors.

[1043] An example system includes an industrial system including an oil refinery. An example oil refinery includes one or more compressors for transferring fluids throughout the plant, and/or for pressurizing fluid streams (e.g., for reflux in a distillation column). Additionally, or alternatively, the example oil refinery includes vacuum distillation, for example to fractionate hydrocarbons. The example oil refinery additionally includes various pipelines in the system for transferring fluids, bringing in feedstock, final product delivery, and the like. An example system includes a number of sensors configured to determine each aspect of a distillation column -- for example temperatures of various fluid streams, temperatures, and compositions of individual contact trays in the column, measurements of the feed and reflux, as well as of the effluent or separated products. The design of a distillation column is complex, and optimal design can depend upon the sizing of boilers, compressors, the contact conditions within the column, as well as the composition of feedstock, which can vary significantly. Additionally, the optimal position for effective sensing of conditions in a pipeline can vary with fluid flow rates, environmental conditions (e.g., causing variation in heat transfer rates), the feedstock utilized, and other factors. Additionally, wear or loss of capability in a boiler, compressor, or other operating equipment can change the system response and capabilities, rendering a single point optimization, including where sensors should be positioned and how they should sample data, to be non-optimal as the system ages.

[1044] Provision of multiple sensors throughout the system can be costly, not necessarily because the sensors are expensive, but because the sensors provide data that may be prohibitive to transmit, store, and utilize. The example system includes providing a large number of sensors throughout the system, and predicting the future states of components, process variables, products, and/or emissions for the system. The example system utilizes a pattern recognition circuit to determine not only the future predicted state of parameters, but when the future predicted state of parameters

will be of interest, and/or will combine with other future predicted state of parameters to create additional risks or opportunities.

[1045] Additionally, the system characterization circuit and the system collaboration circuit can improve predictions and/or system characterizations over time, and/or utilizing offset oil refineries, to more robustly make predictions or system characterizations, which can provide for earlier detection, longer term planning for overall enterprise optimization, and/or to allow the industrial system to operate closer to margins. If an unexpected operating condition occurs -- for example an off-nominal operation of a compressor, the sensor collaboration circuit is able to migrate the system prediction and improve the capability to detect the conditions that caused the unexpected operating condition in the system, and/or in offset systems. Additionally, alerts for the distillation column, based upon predictions indicating off-nominal operation, marginal operation, high risk operation, and/or upcoming maintenance or potential failures, can be readily prepared to provide visibility to risks that otherwise may not be apparent by simply looking at system capacities and past experience without rigorous analysis.

[1046] Example sensor fusion operations for a refinery include vibration information combined with temperatures, pressures, and/or composition (e.g., to determine compressor performance); temperature and pressure, temperature and composition, and/or composition, and pressure (e.g., to determine feedstock variance, contact tray performance, and/or a component failure).

[1047] An example refinery system includes storage tanks and/or boiler feed water. Example system determinations include a sensor fusion to determine a storage tank failure and/or off-nominal operation, such as through a temperature and pressure fusion, and/or a vibration determination with a non-vibration determination (e.g., detecting leaks, air in the system, and/or a feed pump issue). Certain further example system predictions include a sensor fusion to determine a boiler feed water failure, such as through a sensor fusion including flow rate, pressure, temperature, and/or vibration. Any one or more of these parameters can be utilized to predict a system leak, failure, wear of a feed pump, and/or scaling.

[1048] Similarly, an example industrial system includes a power generation system having a condensate and/or make-up water system, where a sensor fusion provides for a sensed parameter group and prediction of failures, maintenance, and the like. The system characterization circuit, utilizing sensor fusion and/or a continuous machine learning process, can predict failures, off-nominal operations, component health, and/or maintenance events for, without limitation, compressors, piping, storage tanks, and/or boiler feed water for an oil refinery.

[1049] An example industrial system includes an irrigation system for a field or a system of fields. Irrigations systems are subject to significant variability in the system (e.g., inlet pressures and/or water levels, component wear and maintenance) as well as environmental variability (e.g., types and distribution of crops planted, weather, soil moisture, humidity, seasonal variability in the sun, cloud coverage, and/or wind variance). Additionally, irrigation systems tend to be remotely located where high bandwidth network access, maintenance facilities, and/or even personnel for oversight are not readily available. An example system includes a multiplicity of sensors capable to enable prediction of conditions for the irrigation system, without requiring that all of the sensors transmit

or store data on a continuous basis. The pattern recognition circuit can readily determine the most important set of sensors to effectively predict patterns and thus system conditions requiring a response (e.g., irrigation cycles, positioning, and the like). Additionally, alerts for remote facilities can be readily prepared, with confidence that the correct sensor package is in place for predicting an off-nominal condition (e.g., imminent failure or maintenance requirement for a pump). In certain example embodiments, the system may determine an off-nominal process condition such as water feed availability being below normal (e.g., based upon recognized pattern conditions such as recent precipitation history, water production history from the irrigation system or other systems competing for the same water feed), structured news alerts or external data, etc., and update the sensed parameter group, for example to confirm the water feed availability (e.g., a water level sensor in a relevant location), to confirm that acceptable conditions are available that water delivery levels can be dropped (e.g., a humidity sensor, and/or a prompt to a user), and/or to confirm that sufficient available secondary sources are available (e.g., an auxiliary water level sensor).

[1050] An example industrial system includes a chemical or pharmaceutical plant. Chemical plants require specific operating conditions, flow rates, temperatures, and the like to maintain proper temperatures, concentrations, mixing, and the like throughout the system. In many systems, there are numerous process steps, and an off-nominal or uncoordinated operation in one part of the process can result in reduced yields, a failed process, and/or a significant reduction in production capacity as coordinated processes must respond (or as coordinated processes fail to respond). Accordingly, a very large number of systems are required to minimally define the system, and in certain embodiments a prohibitive number of sensors are required, from a data transmission and storage viewpoint, to keep sensing capability for a broad range of operating conditions. Additionally, the complexity of the system results in difficulty optimizing and coordinating system operations even where sufficient sensors are present. In certain example embodiments, the pattern recognition circuit can predict the sensing parameter groups that provide high resolution understanding of the system, without requiring that all of the sensors store and transmit data continuously. Further, the pattern recognition circuit can highlight the predicted system risks and capacity limitations for upcoming process operations, where the risks are buried in the complex process. Accordingly, this means it can confidently be operated closer to margins, at a lower cost, and/or maintenance or system upgrades can be performed before failures or capacity limitations are experienced.

[1051] Further, the utilization of a sensor fusion provides for the opportunity to abstract desired predictions, such as "maximize quality" or "minimize and undesirable side reaction" without requiring a full understanding from the operator of which sensors and system conditions are most effective to achieve the abstracted desired output. Further, the predictive nature of the pattern recognition circuit allows for changes in the process to support the desired outcome to be implemented before the process is committed to a sub-optimal outcome. Example components in a chemical or pharmaceutical plan amenable to control and predictions based on operations of the pattern recognition circuit and/or a sensor fusion operation include an agitator, a pressure reactor, a catalytic reactor, and/or a thermic heating system. Example sensor fusion operations to determine

sensed parameter groups and tune the pattern recognition circuit include, without limitation, a vibration sensor combined with another sensor type, a composition sensor combined with another sensor type, a flow rate determination combined with another sensor type, and/or a temperature sensor combined with another sensor type. For example, agitators are amenable to vibration sensing, as well as uniformity of composition detection (e.g., high resolution temperature), expected reaction rates in a properly mixed system, and the like. Catalytic reactors are amenable to temperature sensing (based on the reaction thermodynamics), composition detection (e.g., for expected reactants, as well as direct detection of catalytic material), flow rates (e.g., gross mechanical failure, reduced volume of beads, etc.), and/or pressure detection (e.g., indicative of or coupled with flow rate changes).

[1052] An example industrial system includes a food processing system. Example food processing systems include pressurization vessels, stirrers, mixers, and/or thermic heating systems. Control of the process is critical to maintain food safety, product quality, and product consistency. However, most input parameters to the food processing system are subject to high variability – for example basic food products are inherently variable as natural products, with differing water content, protein content, and other aesthetic variation. Additionally, labor cost management, power cost management, and variability in supply water, etc., provide for a complex process where determination of the predictive variables, sensed parameters to determine these, and optimization of predicting in response to process variation are a difficult problem to resolve. Food processing systems are often cost conscious, and capital costs (e.g., for a robust network and computing system for optimization) are not readily incurred. Further, a food processing system may manufacture wide variance of products on similar or the same production facilities, for example to support an entire product line and/or due to seasonal variations, and accordingly a predictive operation for one process may not support another process well. Example systems include the pattern recognition circuit determining the sensing parameter groups that provide a strong signal response in target outcomes even in light of high variability in system conditions. The pattern recognition circuit can provide for numerous sensed group parameter options available for different process conditions without requiring extensive computing or data storage resources, and accordingly achieve relevant predictions for a wide variety of operating conditions. For example, control of and predictions for pressurization vessels, stirrers, mixers, and/or thermic heating systems are amenable to operations of the pattern recognition circuit, and/or a sensor fusion with a temperature determination combined with a non-temperature determination, a vibration determination combined with a non-vibration determination, and/or a heat map combined with a rate of change in the heat map and/or a non-heat map determination. An example system includes a pattern recognition circuit operation and/or a sensor fusion with a vibration determination and a non-vibration determination, wherein predictive information for a mixer and/or a stirrer is provided; and/or with a pressure determination, a temperature determination, and/or a non-pressure determination, wherein predictive information for a pressurization vessel is provided.

[1053] Referencing Figure 109, an example procedure 11038 includes an operation 11040 to provide a number of sensors to an industrial system including a number of components, each of

the number of sensors operatively coupled to at least one of the number of components, an operation 11042 to interpret a number of sensor data values in response to a sensed parameter group, the sensed parameter group including at least one sensor of the number of sensors, an operation 11044 to determine a recognized pattern value in response to a least a portion of the number of sensor data values, and an operation 11046 to provide a system characterization value for the industrial system in response to the recognized pattern value.

[1054] An example procedure 11038 further includes the operation 11046 to provide the system characterization value by performing an operation such as: determining a predicted outcome for a process associated with the industrial system; determining a predicted future state for a process associated with the industrial system; determining a predicted off-nominal operation for the process associated with the industrial system; determining a prediction value for one of the plurality of components; determining a future state value for one of the plurality of components; determining an anticipated maintenance health state information for one of the plurality of components; determining a predicted maintenance interval for at least one of the plurality of components; determining a predicted off-nominal operation for one of the plurality of components; determining a predicted fault operation for one of the plurality of components; determining a predicted exceedance value for one of the plurality of components; and/or determining a predicted saturation value for one of the plurality of sensors.

[1055] An example procedure 11038 includes an operation 11050 to interpret external data and/or cloud-based data, and where the operation 11044 to determine the recognized pattern value is further in response to the external data and/or the cloud-based data. An example procedure 11038 includes an operation to iteratively improve pattern recognition operations in response to the external data and/or the cloud-based data, for example by operation 11048 to adjust the operation 11042 interpreting sensor values, such as by updating the sensed parameter group. The operation to iteratively improve pattern recognition may further include repeating operations 11042 through 11048, periodically, at selected intervals, in response to a system change, and/or in response to a prediction value of a component, process, or the system.

[1056] In example embodiments, a system for data collection in an industrial environment may comprise: an industrial system comprising a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components; a sensor communication circuit structured to interpret a plurality of sensor data values in response to a sensed parameter group, the sensed parameter group comprising at least one sensor of the plurality of sensors; a pattern recognition circuit structured to determine a recognized pattern value in response to a least a portion of the plurality of sensor data values; and a system characterization circuit structured to provide a system characterization value for the industrial system in response to the recognized pattern value. In example embodiments, a characterization value may include at least one characterization value selected from the characterization values consisting of: a predicted outcome for a process associated with the industrial system; a predicted future state for a process associated with the industrial system; and a predicted off-nominal operation for the process associated with the industrial system. The system characterization value may include at least one characterization

value selected from the characterization values consisting of: a prediction value for one of the plurality of components; a future state value for one of the plurality of components; an anticipated maintenance health state information for one of the plurality of components; and a predicted maintenance interval for at least one of the plurality of components. The system characterization value may include at least one characterization value selected from the characterization values consisting of: a predicted off-nominal operation for one of the plurality of components; a predicted fault operation for one of the plurality of components; and a predicted exceedance value for one of the plurality of components. The system characterization value may include a predicted saturation value for one of the plurality of sensors. A system collaboration circuit may be included that is structured to interpret external data, and wherein the pattern recognition circuit is further structured to determine the recognized pattern value further in response to the external data. The pattern recognition circuit may be further structured to iteratively improve pattern recognition operations in response to the external data. The external data may include at least one of: an indicated component maintenance event; an indicated component failure event; an indicated component wear value; an indicated component operational exceedance value; and an indicated fault value. The external data may include at least one of: an indicated process failure value; an indicated process success value; an indicated process outcome value; and an indicated process operational exceedance value. The external data may include an indicated sensor saturation value. A system collaboration circuit may be included that is structured to interpret cloud-based data comprising a second plurality of sensor data values, the second plurality of sensor data values corresponding to at least one offset industrial system, and wherein the pattern recognition circuit is further structured to determine the recognized pattern value further in response to the cloud-based data. The pattern recognition circuit may be further structured to iteratively improve pattern recognition operations in response to the cloud-based data. The sensed parameter group may include a triaxial vibration sensor. The sensed parameter group may include a vibration sensor and a second sensor that is not a vibration sensor, such as where the second sensor comprises a digital sensor. The sensed parameter group may include a plurality of analog sensors.

[1057] In example embodiments, a method may comprise: providing a plurality of sensors to an industrial system comprising a plurality of components, each of the plurality of sensors operatively coupled to at least one of the plurality of components; interpreting a plurality of sensor data values in response to a sensed parameter group, the sensed parameter group comprising at least one sensor of the plurality of sensors; determining a recognized pattern value in response to a least a portion of the plurality of sensor data values; and providing a system characterization value for the industrial system in response to the recognized pattern value. The system characterization value may be provided by performing at least one operation selected from the operations consisting of: determining a prediction value for one of the plurality of components; determining a future state value for one of the plurality of components; determining an anticipated maintenance health state information for one of the plurality of components; and determining a predicted maintenance interval for at least one of the plurality of components. The system characterization value may be provided by performing at least one operation selected from the operations consisting of:

determining a predicted outcome for a process associated with the industrial system; determining a predicted future state for a process associated with the industrial system; and determining a predicted off-nominal operation for the process associated with the industrial system. The system characterization value may be provided by performing at least one operation selected from the operations consisting of: determining a predicted off-nominal operation for one of the plurality of components; determining a predicted fault operation for one of the plurality of components; and determining a predicted exceedance value for one of the plurality of components. The system characterization value may be provided by determining a predicted saturation value for one of the plurality of sensors. Determining the recognized pattern value may be further in response to the external data. Iteratively improving pattern recognition operations may be provided in response to the external data. Interpreting the external data may include at least one operation selected from the operations consisting of: interpreting an indicated component maintenance event; interpreting an indicated component failure event; interpreting an indicated component wear value; interpreting an indicated component operational exceedance value; and interpreting an indicated fault value. Interpreting the external data may include at least one operation selected from the operations consisting of: interpreting an indicated process success value; interpreting an indicated process failure value; interpreting an indicated process outcome value; and interpreting an indicated process operational exceedance value. Interpreting the external data may include interpreting an indicated sensor saturation value. Interpreting cloud-based data may include a second plurality of sensor data values, the second plurality of sensor data values corresponding to at least one offset industrial system, and wherein determining the recognized pattern value is further in response to the cloud-based data. Iteratively improving pattern recognition operations may be provided in response to the cloud-based data.

[1058] In example embodiments, a system for data collection in an industrial environment may comprise: an industrial system comprising a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components; a sensor communication circuit structured to interpret a plurality of sensor data values in response to a sensed parameter group, the sensed parameter group comprising at least one sensor of the plurality of sensors; a means for determining a recognized pattern value in response to at least a portion of the plurality of sensor data values; and a means for providing a system characterization value for the industrial system in response to the recognized pattern value. The means for providing the system characterization value may comprise a means for performing at least one operation selected from the operations consisting of: determining a predicted outcome for a process associated with the industrial system; determining a predicted future state for a process associated with the industrial system; and determining a predicted off-nominal operation for the process associated with the industrial system. The means for providing the system characterization value may include a means for performing at least one operation selected from the operations consisting of: determining a prediction value for one of the plurality of components; determining a future state value for one of the plurality of components; determining an anticipated maintenance health state information for one of the plurality of components; and determining a predicted maintenance interval for at least

one of the plurality of components. The means for providing the system characterization value may include a means for performing at least one operation selected from the operations consisting of: determining a predicted off-nominal operation for one of the plurality of components; determining a predicted fault operation for one of the plurality of components; and determining a predicted exceedance value for one of the plurality of components. The means for providing the system characterization value may include a means for determining a predicted saturation value for one of the plurality of sensors. A system collaboration circuit may be provided that is structured to interpret external data, and wherein the means for determining the recognized pattern value determines the recognized pattern value further in response to the external data. A means for iteratively improving pattern recognition operations may be provided in response to the external data. The external data may include at least one of: an indicated process success value; an indicated process failure value; and an indicated process outcome value. The external data may include at least one of: an indicated component maintenance event; an indicated component failure event; and an indicated component wear value. The external data may include at least one of: an indicated process operational exceedance value; an indicated component operational exceedance value; and an indicated fault value. The external data may include an indicated sensor saturation value. A system collaboration circuit may be provided that is structured to interpret cloud-based data comprising a second plurality of sensor data values, the second plurality of sensor data values corresponding to at least one offset industrial system, and wherein the means for determining the recognized pattern value determines the recognized pattern value further in response to the cloud-based data. A means for iteratively improving pattern recognition operations may be provided in response to the cloud-based data.

[1059] In example embodiments, a system for data collection in an industrial environment may comprise: a distillation column comprising a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components; a sensor communication circuit structured to interpret a plurality of sensor data values in response to a sensed parameter group, the sensed parameter group comprising at least one sensor of the plurality of sensors; a pattern recognition circuit structured to determine a recognized pattern value in response to a least a portion of the plurality of sensor data values; and a system characterization circuit structured to provide a system characterization value for the distillation column in response to the recognized pattern value. The plurality of components may include a thermodynamic treatment component, and wherein the system characterization value comprises at least one value selected from the values consisting of: determining a prediction value for the thermodynamic treatment component; determining a future state value for the thermodynamic treatment component; determining an anticipated maintenance health state information for the thermodynamic treatment component; and determining a process rate limitation according to a capacity of the thermodynamic treatment component. The thermodynamic treatment component may include at least one of a compressor or a boiler.

[1060] In example embodiments, a system for data collection in an industrial environment may comprise: a chemical process system comprising a plurality of components, and a plurality of

sensors each operatively coupled to at least one of the plurality of components; a sensor communication circuit structured to interpret a plurality of sensor data values in response to a sensed parameter group, the sensed parameter group comprising at least one sensor of the plurality of sensors; a pattern recognition circuit structured to determine a recognized pattern value in response to a least a portion of the plurality of sensor data values; and a system characterization circuit structured to provide a system characterization value for the chemical process system in response to the recognized pattern value. The chemical process system may include one of a chemical plant, a pharmaceutical plant, or an oil refinery. The system characterization value may include at least one value selected from the values consisting of: a separation process value comprising at least one of a capacity value or a purity value; a side reaction process value comprising a side reaction rate value; and a thermodynamic treatment value comprising one of a capability, a capacity, and an anticipated maintenance health for a thermodynamic treatment component.

[1061] In example embodiments, a system for data collection in an industrial environment includes an irrigation system comprising a plurality of components including a pump, and a plurality of sensors each operatively coupled to at least one of the plurality of components; a sensor communication circuit structured to interpret a plurality of sensor data values in response to a sensed parameter group, the sensed parameter group comprising at least one sensor of the plurality of sensors; a pattern recognition circuit structured to determine a recognized pattern value in response to a least a portion of the plurality of sensor data values; and a system characterization circuit structured to provide a system characterization value for the irrigation system in response to the recognized pattern value. The system characterization value may include at least one of an anticipated maintenance health value for the pump and a future state value for the pump. The pattern recognition circuit may determine an off-nominal process condition in response to the at least a portion of the plurality of sensor data values, and wherein the sensor communication circuit is further structured to change the sensed parameter group in response to the off-nominal process condition. The off-nominal process condition may include an indication of below normal water feed availability, and wherein the updated sensed parameter group comprises at least one sensor selected from the sensors consisting of: a water level sensor, a humidity sensor, and an auxiliary water level sensor.

[1062] As described elsewhere herein, feedback to various intelligent and/or expert systems, control systems (including remote and local systems, autonomous systems, and the like), and the like, which may comprise rule-based systems, model-based systems, artificial intelligence (AI) systems (including neural nets, self-organizing systems, and others described throughout this disclosure), and various combinations and hybrids of those (collectively referred to herein as the "expert system" except where context indicates otherwise), may include a wide range of information, including measures such as utilization measures, efficiency measures (e.g., power, financial such as reduction of costs), measures of success in prediction or anticipation of states (e.g., avoidance and mitigation of faults), productivity measures (e.g., workflow), yield measures,

profit measures, and the like, as described herein. In embodiments feedback to the expert system may be industry-specific, domain-specific, factory-specific, machine-specific and the like.

[1063] Industry-specific feedback for the expert system may be offered by a third party, such as a repair and maintenance organization, manufacturer, one or more consortia, and the like, or may be generated by one or more elements of the subject system itself. Industry-specific feedback may be aggregated, such as into one or more data structures, wherein the data are aggregated at the component level, equipment level, factory/installation level, and/or industry level. Users of the data structure(s) may access data at any level (e.g., component, equipment, factory, industry, etc.) Users may search the data structure(s) for indicators/predictors based on or filtered by system conditions specific to their need, or update an indicator/predictor with proprietary data to customize the data structure to their industry. In example embodiments, the expert system may be seeded with industry-specific feedback, such as in a deep learning fashion, to provide an anticipated outcome or state and/or to perform actions to optimize specific machines, devices, components, processes, and the like.

[1064] In example embodiments, feedback provided to the expert system may be used in one or more smart bands to predict progress towards one or more goals. The expert system may use the feedback to determine a modification, alteration, addition, change, or the like to one or more components of the system that provided the feedback, as described elsewhere herein. Based on the industry-specific feedback, the expert system may alter an input, a way of treating or storing an input or output, a sensor or sensors used to provide feedback, an operating parameter, a piece of equipment used in the system, or any other aspect of the participants in the industrial system that gave rise to the feedback. As described elsewhere herein, the expert system may track multiple goals, such as with one or more smart bands. Industry-specific feedback may be used in or by the smart bands in predicting an outcome or state relating to the one or more goals, and to recommend or instruct a change that is directed in increasing a likelihood of achieving the outcome or state.

[1065] For example, a mixer may be used in a food processing environment or in a chemical processing environment, but the feedback that is relevant in the food processing plant (e.g., required sterilization temperatures, food viscosity, particle density (e.g., such as measured by an optical sensor), completion of cooking (e.g., completion of reactions involved in baking), sanitation (e.g., absence of pathogens) may be different than what is relevant in the chemical processing plant (e.g., impeller speed, velocity vectors, flow rate, absence of high contaminant levels, or the like). This industry specific feedback is useful in optimizing the operation of the mixer in its particular environment.

[1066] In another example, the expert system may use feedback from agricultural systems to train a model related to an irrigation system deployed in a field, wherein the industry-specific feedback relates to one or more of an amount of water used across the industry (e.g., such as measured by a flowmeter), a trend of water usage over a time period (e.g., such as measured by a flowmeter), a harvest amount (e.g., such as measured by a weight scale), an insect infestation (e.g., such as identified and/or measured by a drone imaging), a plant death (e.g., such as identified and/or measured by drone imaging), and the like.

[1067] In another example of a fluid flow system (e.g., fan, pump or compressor) controlling cooling in the manufacturing industry, the expert system may use feedback from manufacturing of components involving materials (e.g., polymers) that require cooling during the manufacturing process, such as one or more of quality of output product, strength of output product, flexibility of output product, and the like (e.g., such as measured by a suite of sensors, including densitometer, viscometer, size exclusion chromatograph, and torque meter). If the sensors indicate that the polymer is cooling too quickly during monomer conversion, the expert system may relay an instruction to one or more of a fan, pump, or compressor in the fluid flow system to decrease an aspect of its operation in order to meet a quality goal.

[1068] In another example of a reciprocating compressor operating in a refinery performing refinery processes (e.g., hydrotreating, hydrocracking, isomerization, reforming), the expert system may use feedback related to one or more of an amount of sulfur, nitrogen and/or aromatics downstream of the compressor (e.g., such as measured by a near infrared ("IR") analyzer), the cetane/octane number or smoke point of a product (e.g., such as with an octane analyzer), the density of a product (e.g., such as measured by a densitometer), byproduct gas amounts (e.g., such as measured by an electrochemical gas sensor), and the like. In this example, as feedback is received during isomerization of butane to isobutene by an inline near IR analyzer measuring the amount and/or quality of isobutene, the expert system may determine that the performance of one or more components of the isomerization system, including the reciprocating compressor, should be altered in order to meet a production goal.

[1069] In another example of a vacuum distillation unit operating in a refinery, the expert system may use feedback related to an amount of raw gasoline recovered (e.g., such as by measuring the volume or composition of various fractions using IR), boiling point of recovered fractions (e.g., such as with a boiling point analyzer), a vapor cooling rate (e.g., such as measured by thermometer), and the like. In this example, as feedback is received during vacuum distillation to recover diesel, as the amounts recovered indicate off-nominal rations of production, the expert system may instruct the vacuum distillation unit to alter a feedstock source and initiate more detailed analysis of the prior feedstock.

[1070] In yet another example of a pipeline in a refinery, the expert system may use feedback related to flow type (e.g., bubble, stratified, slug, annular, transition, mist) of hydrocarbon products (e.g., such as measured by dye tracing), flow rate, vapor velocity (such as with a flow meter), vapor shear, and the like. In this example, as feedback is received during operation of the pipeline regarding the flow type and its rate, modifications may be recommended by the expert system to improve the flow through the pipeline.

[1071] In still another example of a paddle-type or anchor-type agitator/mixer in a pharmaceutical plant, the expert system may use feedback related to degree of mixing of high-viscosity liquids, heating of medium- to low-viscosity liquids, a density of the mixture, a growth rate of an organism in the mixture, and the like. In this example, as feedback is received during operation of the agitator that a bacterial growth rate is too high (such as measured with a spectrophotometer), the expert

system may instruct the agitator to reduce its speed to limit the amount of air being added to the mixture or growth substrate.

[1072] In a further example of a pressure reactor in a chemical processing plant, the expert system may use feedback related to a catalytic reaction rate (such as measured by a mass spectrometer), a particle density (such as measured by a densitometer), a biological growth rate (such as measured by a spectrophotometer), and the like. In this example, as feedback is received during operation of the pressure reactor that the particle density and biological growth rate are off-nominal, the expert system may instruct the pressure reactor to modify one or more operational parameters, such as a reduction in pressure, an increase in temperature, an increase in volume of the reaction, and the like.

[1073] In another example of a gas agitator operating in a chemical processing plant, the expert system may use feedback related to effective density of a gassed liquid, a viscosity, a gas pressure, and the like, as measured by appropriate sensors or equipment. In this example, as feedback is received during operation of the gas agitator, the expert system may instruct the gas agitator to modify one or more operational parameters, such as to increase or decrease a rate of agitation.

[1074] In still another example of a pump blasting liquid type agitator in a chemical processing plant, the expert system may use feedback related to a viscosity of a mixture, an optical density of a growth medium, and a temperature of a solution. In this example, as feedback is received during operation of the agitator, the expert system may instruct the agitator to modify one or more operational parameters, such as to increase or decrease a rate of agitation and/or inject additional heat.

[1075] In yet another example of a turbine type agitator in a chemical processing plant, the expert system may use feedback related to a vibration noise, a reaction rate of the reactants, a heat transfer, or a density of a suspension. In this example, as feedback is received during operation of the agitator, the expert system may instruct the agitator to modify one or more operational parameters, such as to increase or decrease a rate of agitation and/or inject an additional amount of catalyst.

[1076] In yet another example of a static agitator mixing monomers in a chemical processing plant to produce a polymer, the expert system may use feedback related to the viscosity of the polymer, color of the polymer, reactivity of the polymer and the like to iterate to a new setting or parameter for the agitator, such as for example, a setting that alters the Reynolds number, an increase in temperature, a pressure increase, and the like.

[1077] In a further example of a catalytic reactor in a chemical processing plant, the expert system may use feedback related to a reaction rate, a product concentration, a product color, and the like. In this example, as feedback is received during operation of the catalytic reactor, the expert system may instruct the reactor to modify one or more operational parameters, such as to increase or decrease a temperature and/or inject an additional amount of catalyst.

[1078] In yet a further example of a thermic heating system in a chemical processing or food plant, the expert system may use feedback related to BTUs out of the system, a flow rate, and the like. In this example, as feedback is received during operation of the thermic heating system, the expert

system may instruct the system to modify one or more operational parameters, such as to change the input feedstock, to increase the flow of the feedstock, and the like.

[1079] In still a further example of using boiler feed water in a refinery, the expert system may use feedback related to an aeration level, a temperature, and the like. In this example, as feedback is received related to the boiler feed water, the expert system may instruct the system to modify one or more operational parameters of a boiler, such as to increase a reduction in aeration, to increase the flow of the feed water, and the like.

[1080] In still a further example of a storage tank in a refinery, the expert system may use feedback related to a temperature, a pressure, a flow rate out of the tank, and the like. In this example, as feedback is received related to the storage tank, the expert system may instruct the system to modify one or more operational parameters of, such as to increase cooling or heating begin agitation, and the like.

[1081] In an example of a condensate/make-up water system in a power station that condenses steam from turbines and recirculates it back to a boiler feeder along with make-up water, the expert system may use feedback related to measuring inward air leaks, heat transfer, and make-up water quality. In this example, as feedback is received related to the condensate/make-up water system, the expert system may instruct the system to increase a purification of the make-up water, bring a vacuum pump online, and the like.

[1082] In another example of a stirrer in a food plant, the expert system may use feedback related to a viscosity of the food, a color of the food, a temperature of the food, and the like. In this example, as feedback is received, the expert system may instruct the stirrer to speed up or slow down, depending on the predicted success in reaching a goal.

[1083] In another example of a pressure cooker in a food plant, the expert system may use feedback related to a viscosity of the food, a color of the food, a temperature of the food, and the like. In this example, as feedback is received, the expert system may instruct the pressure cooker to continue operating, increase a temperature, or the like, depending on the predicted success in reaching a goal.

[1084] In example embodiments, a system 11100 for data collection in an industrial environment may include a plurality of input sensors 11102 communicatively coupled to a controller 11106, a data collection circuit 11104 structured to collect output data 11108 from the input sensors 11102, and a machine learning data analysis circuit 11110 structured to receive the output data 11108 and learn received output data patterns 11112 indicative of an outcome, wherein the machine learning data analysis circuit 11110 is structured to learn received output data patterns 11112 by being seeded with a model 11114 based on industry-specific feedback 11118. The model 11114 may be a physical model, an operational model, or a system model. The industry-specific feedback 11118 may be one or more of a utilization measure, an efficiency measure (e.g., power and/or financial), a measure of success in prediction or anticipation of states (e.g., an avoidance and mitigation of faults), a productivity measure (e.g., a workflow), a yield measure, and a profit measure. The industry-specific feedback 11118 includes an amount of power generated by a machine about which the input sensors provide information during operation of the machine. The industry-specific

feedback 11118 includes a measure of the output of an assembly line about which the input sensors provide information. The industry-specific feedback 11118 includes a failure rate of units of product produced by a machine about which the input sensors provide information. The industry-specific feedback 11118 includes a fault rate of a machine about which the input sensors provide information. The industry-specific feedback 11118 includes the power utilization efficiency of a machine about which the input sensors provide information, wherein the machine is one of a turbine, a transformer, a generator, a compressor, one that stores energy, and one that includes power train components (e.g., the rate of extraction of a material by a machine about which the input sensors provide information, the rate of production of a gas by a machine about which the input sensors provide information, the rate of production of a hydrocarbon product by a machine about which the input sensors provide information), and the rate of production of a chemical product by a machine about which the input sensors provide information. The machine learning data analysis circuit 11110 may be further structured to learn received output data patterns 11112 based on the outcome. The system 11100 may keep or modify operational parameters or equipment. The controller 11106 may adjust the weighting of the machine learning data analysis circuit 11110 based on the learned received output data patterns 11112 or the outcome, collect more/fewer data points from the input sensors based on the learned received output data patterns 11112 or the outcome, change a data storage technique for the output data 11108 based on the learned received output data patterns 11112 or the outcome, change a data presentation mode or manner based on the learned received output data patterns 11112 or the outcome, and apply one or more filters (low pass, high pass, band pass, etc.) to the output data 11108. In example embodiments, the system 11100 may remove/re-task under-utilized equipment based on one or more of the learned received output data patterns 11112 and the outcome. The machine learning data analysis circuit 11110 may include a neural network expert system. The input sensors may measure vibration and noise data. The machine learning data analysis circuit 11110 may be structured to learn received output data patterns 11112 indicative of progress/alignment with one or more goals/guidelines (e.g., which may be determined by a different subset of the input sensors). The machine learning data analysis circuit 11110 may be structured to learn received output data patterns 11112 indicative of an unknown variable. The machine learning data analysis circuit 11110 may be structured to learn received output data patterns 11112 indicative of a preferred input among available inputs. The machine learning data analysis circuit 11110 may be structured to learn received output data patterns 11112 indicative of a preferred input data collection band among available input data collection bands. The machine learning data analysis circuit 11110 may be disposed in part on a machine, on one or more data collectors, in network infrastructure, in the cloud, or any combination thereof. The system 11100 may be deployed on the data collection circuit 11104. The system 11100 may be distributed between the data collection circuit 11104 and a remote infrastructure. The data collection circuit 11104 may include a data collector.

[1085] In example embodiments, a system 11100 for data collection in an industrial environment may include a plurality of input sensors 11102 communicatively coupled to a controller 11106, a data collection circuit 11104 structured to collect output data 11108 from the input sensors, and a

machine learning data analysis circuit 11110 structured to receive the output data 11108 and learn received output data patterns 11112 indicative of an outcome, wherein the machine learning data analysis circuit 11110 is structured to learn received output data patterns 11112 by being seeded with a model 11114 based on a utilization measure.

[1086] In example embodiments, a system 11100 for data collection in an industrial environment may include a plurality of input sensors 11102 communicatively coupled to a controller 11106, a data collection circuit 11104 structured to collect output data 11108 from the input sensors, and a machine learning data analysis circuit 11110 structured to receive the output data 11108 and learn received output data patterns 11112 indicative of an outcome, wherein the machine learning data analysis circuit 11110 is structured to learn received output data patterns 11112 by being seeded with a model 11114 based on an efficiency measure.

[1087] In example embodiments, a system 11100 for data collection in an industrial environment may include a plurality of input sensors 11102 communicatively coupled to a controller 11106, a data collection circuit 11104 structured to collect output data 11108 from the input sensors, and a machine learning data analysis circuit 11110 structured to receive the output data 11108 and learn received output data patterns 11112 indicative of an outcome, wherein the machine learning data analysis circuit 11110 is structured to learn received output data patterns 11112 by being seeded with a model 11114 based on a measure of success in prediction or anticipation of states.

[1088] In example embodiments, a system 11100 for data collection in an industrial environment may include a plurality of input sensors 11102 communicatively coupled to a controller 11106, a data collection circuit 11104 structured to collect output data 11108 from the input sensors, and a machine learning data analysis circuit 11110 structured to receive the output data 11108 and learn received output data patterns 11112 indicative of an outcome, wherein the machine learning data analysis circuit 11110 is structured to learn received output data patterns 11112 by being seeded with a model 11114 based on a productivity measure.

[1089] In example embodiments, a system for data collection in an industrial environment may include an expert system graphical user interface in which a user may, by interacting with a graphical user interface element, set a parameter of a data collection band for collection by a data collector. The parameter may relate to at least one of setting a frequency range for collection and setting an extent of granularity for collection.

[1090] In example embodiments, a system for data collection in an industrial environment may include an expert system graphical user interface in which a user may, by interacting with a graphical user interface element, identify a set of sensors among a larger set of available sensors for collection by a data collector. The user interface may include views of available data collectors, their capabilities, one or more corresponding smart bands, and the like.

[1091] In example embodiments, a system for data collection in an industrial environment may include an expert system graphical user interface in which a user may, by interacting with a graphical user interface element, select a set of inputs to be multiplexed among a set of available inputs.

[1092] In example embodiments, a system for data collection in an industrial environment may include an expert system graphical user interface in which a user may, by interacting with a graphical user interface element, select a component of an industrial machine displayed in the graphical user interface for data collection, view a set of sensors that are available to provide data about the industrial machine, and select a subset of sensors for data collection.

[1093] In example embodiments, a system for data collection in an industrial environment may include an expert system graphical user interface in which a user may, by interacting with a graphical user interface element, view a set of indicators of fault conditions of one or more industrial machines, where the fault conditions are identified by application of an expert system to data collected from a set of data collectors. In example embodiments, the fault conditions may be identified by manufacturers of portions of the one or more industrial machines. The fault conditions may be identified by analysis of industry trade data, third-party testing agency data, industry standards, and the like. In example embodiments, a set of indicators of fault conditions of one or more industrial machines may include indicators of stress, vibration, heat, wear, ultrasonic signature, operational deflection shape, and the like, optionally including any of the widely varying conditions that can be sensed by the types of sensors described throughout this disclosure and the documents incorporated herein by reference.

[1094] In example embodiments, a system for data collection in an industrial environment may include an expert graphical user interface that enables a user to select from a list of component parts of an industrial machine for establishing smart-band monitoring and in response thereto presents the user with at least one smart-band definition of an acceptable range of values for at least one sensor of the industrial machine and a list of correlated sensors from which data will be gathered and analyzed when an out of acceptable range condition is detected from the at least one sensor.

[1095] In example embodiments, a system for data collection in an industrial environment may include an expert graphical user interface that enables a user to select from a list of conditions of an industrial machine for establishing smart-band monitoring and, in response thereto, presents the user with at least one smart-band definition of an acceptable range of values for at least one sensor of the industrial machine and a list of correlated sensors from which data will be gathered and analyzed when an out of acceptable range condition is detected from the at least one sensor.

[1096] In example embodiments, a system for data collection in an industrial environment may include an expert graphical user interface that enables a user to select from a list of reliability measures of an industrial machine for establishing smart-band monitoring and, in response thereto, presents the user with at least one smart-band definition of an acceptable range of values for at least one sensor of the industrial machine and a list of correlated sensors from which data will be gathered and analyzed when an out of acceptable range condition is detected from the at least one sensor. In the system, the reliability measures may include one or more of industry average data, manufacturer's specifications, material specifications, recommendations, and the like. In example embodiments, reliability measures may include measures that correlate to failures, such as stress, vibration, heat, wear, ultrasonic signature, operational deflection shape effect, and the like. In

example embodiments, manufacturer's specifications may include cycle count, working time, maintenance recommendations, maintenance schedules, operational limits, material limits, warranty terms, and the like. In example embodiments, the sensors in the industrial environment may be correlated to manufacturer's specifications by associating a condition being sensed by the sensor to a specification type. In example embodiments, a non-limiting example of correlating a sensor to a manufacturer's specification may include a duty cycle specification being correlated to a sensor that detects revolutions of a moving part. In example embodiments, a temperature specification may correlate to a thermal sensor disposed to sense an ambient temperature proximal to the industrial machine.

[1097] In example embodiments, a system for data collection in an industrial environment may include an expert graphical user interface that automatically creates a smart-band group of sensors disposed in the industrial environment in response to receiving a condition of the industrial environment for monitoring and an acceptable range of values for the condition.

[1098] In example embodiments, a system for data collection in an industrial environment may include an expert graphical user interface that presents a representation of components of an industrial machine deployable in the industrial environment on an electronic display, and in response to a user selecting one or more of the components, searches a database of industrial machine failure modes for modes involving the selected component(s) and conditions associated with the failure mode(s) to be monitored, and further identifies a plurality of sensors in, on, or available to be disposed on the presented machine representation from which data will automatically be captured when the condition(s) to be monitored are detected to be outside of an acceptable range. In example embodiments, the identified plurality of sensors includes at least one sensor through which the condition(s) will be monitored.

[1099] In example embodiments, a system for data collection in an industrial environment may include a user interface for working with smart bands that may facilitate a user identifying sensors to include in a smart band group of sensors by selecting sensors presented on a map of a machine in the environment. A user may be guided to select among a subset of all possible sensors based on smart band criteria, such as types of sensor data required for the smart band. A smart band may be focused on detecting trending conditions in a portion of the industrial environment; therefore, the user interface may direct the user choose among an identified subset of sensors, such as by only allowing sensors proximal to the smart band directed portion of the environment to be selectable in the user interface.

[1100] In example embodiments, a smart band data collection configuration and deployment user interface may include views of components in an industrial environment and related available sensors. In example embodiments, in response to selection of a component part of an industrial machine depicted in the user interface, sensors associated with smart band data collection for the component part may be highlighted so that the user may select one or more of the sensors. User selection in this context may comprise a verification of an automatic selection of sensors, or manually identifying sensors to include in the smart band sensor group.

[1101] In example embodiments, in response to selection of a smart band condition, such as trending of bearing temperature, a user interface for smart band configuration and use may automatically identify and present sensors that contribute to smart band analysis for the condition. A user may respond to this presentation of sensors, confirm or otherwise acknowledge one or more sensors individually or as a set to be included in the smart band data collection group.

[1102] In example embodiments, a smart band user interface may present locations of industrial machines in an industrial environment on a map. The locations may be annotated with indicators of smart band data collection templates that are configured for collecting smart band data for the machines at the annotated locations. The locations may be color coded to reflect a degree of smart band coverage for a machine at the location. In example embodiments, a location of a machine with a high degree of smart band coverage may be colored green, whereas a location of a machine with low smart band coverage may be colored red or some other contrasting color. Other annotations, such as visual annotations may be used. A user may select a machine at a location and by dragging the selected machine to a location of a second machine, effectively configure smart bands for the second machine that correspond to smart bands for the first machine. In this way, a user may configure several smart band data collection templates for a newly added machine or a new industrial environment and the like.

[1103] In example embodiments, various configurations and selections of smart bands may be stored for use throughout a data collection platform, such as for selecting templates for sensing, templates for routing, provisioning of devices and the like, as well as for direct the placement of sensors, such as by personnel or by machines, such as autonomous or remote-control drones.

[1104] In example embodiments, a smart band user interface may present a map of an industrial environment that may include industrial machines, machine-specific data collectors, mobile data collectors (robotic and human), and the like. A user may view a list of smart band data collection actions to be performed and may select a data collection resource set to undertake the collection. In an example, a guided mobile robot may be equipped with data collection systems for collecting data for a plurality of smart band data sets. A user may view an industrial environment with which the robot is associated and assign the robot to perform a smart band data collection activity by selecting the robot, a smart band data collection template, and a location in the industrial environment, such as a machine or a part of a machine. The user interface may provide a status of the collection undertaking so that the user can be informed when the data collection is complete.

[1105] In example embodiments, a smart band operation management user interface may include presentation of smart band data collection activity, analysis of results, actions taken based on results, suggestions for changes to smart band data collection (e.g., addition of sensors to a smart band collection template, increasing duration of data collection for a template-specific collection activity), and the like. The user interface may facilitate "what if" type analysis by presenting potential impacts on reliability, costs, resource utilization, data collection tradeoffs, maintenance schedule impacts, risk of failure (increase/decrease), and the like in response to a user's attempt to make a change to a smart band data collection template, such as a user relaxing a threshold for performing smart band data collection and the like. In example embodiments, a user may select or

enter a target budget for preventive maintenance per unit time (e.g., per month, quarter, and the like) into the user interface and an expert system of the user interface may recommend a smart band data collection template and thresholds for complying with the budget.

[1106] In example embodiments, a smart band user interface may facilitate a user configuring a system for data collection in an industrial environment for smart band data gathering. The user interface may include display of industrial machine components, such as motors, linkages, bearings, and the like that a user may select. In response to such a selection, an expert system may work with the user interface to present a list of potential failure conditions related to the part to monitor. The user may select one or more conditions to monitor. The user interface may present the conditions to monitor as a set that the user may be asked to approve. The user may indicate acceptance of the set or of select conditions in the set monitor. As a follow-on to a user selection/approval of one or more conditions to monitor, the user interface may display a map of relevant sensors available in the industrial environment for collecting data as a smart band group of sensors. The relevant sensors may be associated with one or more parts (e.g., the part(s) originally selected by the user), one or more failure conditions, and the like.

[1107] In example embodiments, the expert system may compare the relevant sensors in the environment to a preferred set of sensors for smart band monitoring of the failure condition(s) and provide feedback to the user, such as a confidence factor for performing smart band monitoring based on the available sensors for the failure condition(s). The user may evaluate the failure condition and smart band analysis information presented and may take an action in the user interface, such as approving the relevant sensors. In response, a smart band data collection template for configuring the data collection system may be created. In example embodiments, a smart band data collection template may be created independently of a user approval. In such embodiments, the user may indicate explicitly or implicitly via approval of the smart band analysis information and approval of the created template.

[1108] In example embodiments, a smart band user interface may work with an expert system to present candidate portions of an industrial machine in an industrial environment for smart band condition monitoring based on information such as manufacturer's specifications, statistical information derived from real-world experience with similar industrial machines, and the like. In example embodiments, the user interface may permit a user to select certain aspects of the smart band data collection and analysis process including – for example, a degree of reliability/failure risk to monitor (e.g., near failure, best performance, industry average, and the like). In response thereto, the expert system may adjust an aspect of the smart band analysis, such as a range of acceptable value to monitor, a monitor frequency, a data collection frequency, a data collection amount, a priority for the data collection activity (e.g., effectively a priority of a template for data collection for the smart band), weightings of data from sensors (e.g., specific sensors in the group, types of sensors, and the like).

[1109] In example embodiments, a smart bands user interface may be structured to allow a user to let an expert system recommend one or more smart bands to implement based on a range of comparative data that the user might prioritize, such as industry average data, industry best data,

near-by comparable machines, most similarly configured machines, and the like. Based on the comparative data weighting, the expert system may use the user interface to recommend one or more smart band templates that align with the weighting to the user, who may take an action in the user interface, such as approving one or more of the recommended templates for use.

[1110] In example embodiments, a user interface for configuring arrangement of sensors in an industrial environment may include recommendations by industrial environment equipment suppliers (e.g., manufacturers, wholesalers, distributors, dealers, third-party consultants, and the like) of group(s) of sensors to include for performing smart band analysis of components of the industrial equipment. The information may be presented to a user as data collection template(s) that the user may indicate as being accepted/approved, such as by positioning a graphic representing a template(s) over a portion of the industrial equipment.

[1111] In example embodiments, a smart band discovery portal may facilitate sharing of smart band related information, such as recommendations, actual use cases, results of smart band data collection and processing, and the like. The discovery portal may be embodied as a panel in a smart band user interface.

[1112] In example embodiments, a smart band assessment portal may facilitate assessment of smart band-based data collection and analysis. Content that may be presented in such a portal may include depictions of uses of existing smart band templates for one or more industrial machines, industrial environments, industries, and the like. A value of a smart band may be ascribed to each smart band in the portal based, for example, on historical use and outcomes. A smart band assessment portal may also include visualization of candidate sensors to include in a smart band data collection template based on a range of factors including ascribed value, preventive maintenance costs, failure condition being monitored, and the like.

[1113] In example embodiments, a smart bands graphical user interface associated with a system for data collection in an industrial environment may be deployed for industrial components, such as of factory-based air conditioning units. A user interface of a system for data collection for smart band analysis of air conditioning units may facilitate graphical configuration of smart band data collection templates and the like for specific air conditioning system installations. In example embodiments, major components of an air conditioning system, such as a compressor, condenser, heat exchanger, ducting, coolant regulators, filters, fans, and the like along with corresponding sensors for a particular installation of the air conditioning system may be depicted in a user interface. A user may select one or more of these components in the user interface for configuring a system for smart band data collection. In response to the user selecting, for example, a coolant compressor, sensors associated with the compressor may be automatically identified in the user interface. The user may be presented with a recommended data collection template to perform smart band data collection for the selected compressor. Alternatively, the user may request a candidate collection template from a community of smart band users, such as through a smart band template sharing panel of the user interface. Once a template is selected, the user interface may offer the user customization options, such as frequency of collection, degree of reliability to monitor, and the like. Upon final acceptance of the template, the user interface may interact with a

data collection system of the installed air conditioning system (if such a system is available) to implement the data collection template and provide an indication to the user of the result of implementing the template. In response thereto, the user may make a final approval of the template for use with the air conditioning unit.

[1114] In example embodiments, a smart bands graphical user interface associated with a system for data collection in an industrial environment may be deployed for oil and gas refinery-based chillers. A user interface of a system for data collection for smart band analysis of refinery-based chillers may facilitate graphical configuration of smart band data collection templates and the like for specific refinery-based chiller installations. In example embodiments, major components of a refinery-based chiller including heat exchangers, compressors, water regulators and the like along with corresponding sensors for the particular installation of the refinery-based chiller may be depicted in a user interface. A user may select one or more of these components in the user interface for configuring a system for smart band data collection. In response to the user selecting, for example, water regulators, sensors associated with the water regulators may be automatically identified in the user interface. The user may be presented with a recommended data collection template to perform smart band data collection for the selected component. Alternatively, the user may request a candidate collection template from a community of smart band users, such as through a smart band template sharing panel of the user interface. Once a template is selected, the user interface may offer the user customization options, such as frequency of collection, degree of reliability to monitor, and the like. Upon final acceptance of the template, the user interface may interact with a data collection system of the installed refinery-based chiller (if such a system is available) to implement the data collection template and provide an indication to the user of the result of implementing the template. In response thereto, the user may make a final approval of the template for use with the refinery-based chiller.

[1115] In example embodiments, a smart bands graphical user interface associated with a system for data collection in an industrial environment may be deployed for automotive production line robotic assembly systems. A user interface of a system for data collection for smart band analysis of production line robotic assembly systems may facilitate graphical configuration of smart band data collection templates and the like for specific production line robotic assembly system installations. In example embodiments, major components of a production line robotic assembly system including motors, linkages, tool handlers, positioning systems and the like along with corresponding sensors for the particular installation of the production line robotic assembly system may be depicted in a user interface. A user may select one or more of these components in the user interface for configuring a system for smart band data collection. In response to the user selecting, for example, robotic linkage sensors associated with the robotic linkages may be automatically identified in the user interface. The user may be presented with a recommended data collection template to perform smart band data collection for the selected component. Alternatively, the user may request a candidate collection template from a community of smart band users, such as through a smart band template sharing panel of the user interface. Once a template is selected, the user interface may offer the user customization options, such as frequency of collection, degree of

reliability to monitor, and the like. Upon final acceptance of the template, the user interface may interact with a data collection system of the installed production line robotic assembly system (if such a system is available) to implement the data collection template and provide an indication to the user of the result of implementing the template. In response thereto, the user may make a final approval of the template for use with the production line robotic assembly system.

[1116] In example embodiments, a smart bands graphical user interface associated with a system for data collection in an industrial environment may be deployed for automotive production line robotic assembly systems. A user interface of a system for data collection for smart band analysis of production line robotic assembly systems may facilitate graphical configuration of smart band data collection templates and the like for specific production line robotic assembly system installations. In example embodiments, major components of construction site boring machinery, such as the cutter head, which itself is a subsystem that may have many components, control systems, debris handling and conveying components, precast concrete delivery and installation subsystems and the like along with corresponding sensors for the particular installation of the production line robotic assembly system may be depicted in a user interface. A user may select one or more of these components in the user interface for configuring a system for smart band data collection. In response to the user selecting, for example, debris handling components sensors associated with the debris handling components, such as a conveyer may be automatically identified in the user interface. The user may be presented with a recommended data collection template to perform smart band data collection for the selected component. Alternatively, the user may request a candidate collection template from a community of smart band users, such as through a smart band template sharing panel of the user interface. Once a template is selected, the user interface may offer the user customization options, such as frequency of collection, degree of reliability to monitor, and the like. Upon final acceptance of the template, the user interface may interact with a data collection system of the installed production line robotic assembly system (if such a system is available) to implement the data collection template and provide an indication to the user of the result of implementing the template. In response thereto, the user may make a final approval of the template for use with the production line robotic assembly system.

[1117] Referring to Figure 111, an exemplary user interface for smart band configuration of a system for data collection in an industrial environment is depicted. The user interface 11200 may include an industrial environment visualization portion 11202 in which may be depicted one or more sensors, machines, and the like. Each sensor, machine, or portion thereof (e.g., motor, compressor, and the like) may be selectable as part of a smart-band configuration process. Likewise, each sensor, machine or portion thereof may be visually highlighted during the smart-band configuration process, such as in response to user selection, or automated identification as being part of a group of smart band sensors. The user interface may also include a smart band selection portion 11204 or panel in which smart band indicators, failure modes, and the like may be depicted in selectable elements. User selection of a symptom, failure mode and the like may cause corresponding components, sensors, machines, and the like in the industrial visualization portion to be highlighted. The user interface may also include a customization panel 11206 in

which attributes of a selected smart band, such as acceptable ranges, frequency of monitoring and the like may be made available for a user to adjust.

[1118] Back calculation, such as for determining possible root causes of failures and the like, may benefit from a graphical approach that facilitates visualizing an industrial environment, machine, or portion thereof marked with indications of information sources that may provide data such as sensors and the like related to the failure. A failed part, such as a bearing, may be associated with other parts, such as shaft, motor, and the like. Sensors for monitoring conditions of the bearing and the associated parts may provide information that could indicate a potential source of failure. Such information may also be useful to suggest indicators, such as changes in sensor output, to monitor or avoid the failure in the future. A system that facilitates a graphical approach for back-calculation may interact with sensor data collection and analysis systems to at least partially automate aspects related to data collection and processing determined from a back-calculation process.

[1119] In example embodiments, a system for data collection in an industrial environment may include a user interface in which portions of an industrial machine associated with a condition of interest, such as a failure condition, are presented on an electronic display along with sensor data types contributing to the condition of interest, data collection points (e.g., sensors) associated with the machine portions that monitor the data types, a set of data from the data collection points that was collected and used to determine the condition of interest, and an annotation of sensors that delivered exceptional data, such as data that is out of an acceptable range, and the like, that may have been used to determine the condition of interest. The user interface may access a description of the machine that facilitates determining and visualizing related components, such as bearings, shafts, brakes, rotors, motor housings, and the like that contribute to a function, such as rotating a turbine. The user interface may also access a data set that relates sensors disposed in and about the machine with the components. Information in the data set may include descriptions of the sensors, their function, a condition that each senses, typical or acceptable ranges of values output from the sensors, and the like. The information in the data set may also identify a plurality of potential pathways in a system for data collection in an industrial environment for sensor data to be delivered to a data collector. The user interface may also access a data set that may include data collection templates used to configure a data collection system for collecting data from the sensors to meet specific purposes (e.g., to collect data from groups of sensors into a sensor data set suitable for determining a condition of the machine, such as a degree of slippage of the shaft relative to the motor, and the like).

[1120] In example embodiments, a method of back-calculation for determining candidate sources of data collection for data that contributes to a condition of an industrial machine may include following routes of data collection determined from a configuration and operational template of a data collection system for collecting data from sensors deployed in the industrial machine that was in place when the contributing data was collected. A configuration and operational template may describe signal path switching, multiplexing, collection timing, and the like for data from a group of sensors. The group of sensors may be local to a component, such as a bearing, or more regionally distributed, such as sensors that capture information about the bearing and its related components.

In example embodiments, a data collection template may be configured for collecting and processing data to detect a particular condition of the industrial machine. Therefore, templates may be correlated to conditions so that performing back-calculation of a condition of interest can be guided by the correlated template. Data collected based on the template may be examined and compared to acceptable ranges of data for various sensors. Data that is outside of an acceptable range may indicate potential root causes of an unacceptable condition. In example embodiments, a suspect source of data collection may be determined from the candidate sources of data collection based on a comparison of data collected from the candidate data sources with an acceptable range of data collected from each candidate data source. Visualizing these back-calculation based signal paths, candidate sensors, and suspect data sources provides a user with valuable insights into possible root causes of failures and the like.

[1121] In example embodiments, a method for back-calculation may include visualizing route(s) of data that contribute to a fault condition detected in an industrial environment by applying back-calculation to determine sources of the contributed data with the visualizing appearing as highlighted data paths in a visual representation of the data collection system in the industrial machine. In example embodiments, determining sources of data may be based on a data collection and processing template for the fault condition. The template may include a configuration of a data collection system when data from the determined sources was collected with the system.

[1122] When failures occur, or conditions of a portion of a machine in an industrial environment reach a critical point prior to failure, such as may be detected during preventive maintenance and the like, back-calculation may be useful in determining information to gather that might help avoid the failure and/or improve system performance – for example, by avoiding substantive degradation in component operation. Visualizing data collection sources, components related to a condition, algorithms that may determine the potential onset of the condition and the like may facilitate preparation of data collection templates for configuring data sensing, routing, and collection resources in a system for data collection in an industrial environment. In example embodiments, configuring a data collection template for a system for collecting data in an industrial environment may be based on back-calculations applied to machine failures that identify candidate conditions to monitor for avoiding the machine failures. The resulting template may identify sensors to monitor, sensor data collection path configuration, frequency, and amount of data to collect, acceptable levels of sensor data, and the like. With access to information about the machine, such as which parts closely relate to others and sensors that collected data from parts in the machine, a data collection system configuration template may be automatically generated when a target component is identified.

[1123] In example embodiments, a user interface may include a graphical display of data sources as a logical arrangement of sensors that may contribute data to a calculation of a condition of a machine in an industrial environment. A logical arrangement may be based on sensor type, data collection template, condition, algorithm for determining a condition, and the like. In an example, a user may wish to view all temperature sensors that may contribute to a condition, such as a failure of a part in an industrial environment. A user interface may communicate with a database of

machine related information, such as parts that relate to a condition, sensors for those parts, and types of those sensors to determine the subset of sensors that measure temperature. The user interface may highlight those sensors. The user interface may activate selectable graphical elements for those sensors that, when selected by the user, may present data associated with those sensors, such as sensor type, ranges of data collected, acceptable ranges, actual data values collected for a given condition, and the like, such as in a pop-up panel or the like. Similar functionality of the user interface may apply to physical arrangements of sensors, such as all sensors associated with a motor, boring machine cutting head, wind turbine, and the like.

[1124] In example embodiments, third-parties, such as component manufacturers, remote maintenance organizations and the like may benefit from access to back-calculation visualization. Permitting third parties to have access to back-calculation information, such as sensors that contributed unacceptable data values to a calculation of a condition, visualization of sensor positioning, and the like may be an option that a user can exercise in a user interface for a graphical approach to back-calculations as described herein. A list of manufacturers of machines, sub-systems, individual components, sensors, data collection systems, and the like may be presented along with remote maintenance organizations, and the like in a portion of a user interface. A user of the interface may select one or more of these third-parties to grant access to at least a portion of the available data and visualizations. Selecting one or more of these third-parties may also present statistical information about the party, such occurrences and frequency of access to data to which the party is granted access, request from the party for access, and the like.

[1125] In example embodiments, visualization of back-calculation analysis may be combined with machine learning so that back-calculations and their visualizations may be used to learn potential new diagnoses for conditions, such as failure conditions, to learn new conditions to monitor, and the like. A user may interact with the user interface to provide the machine learning techniques feedback to improve results, such as indicating a success or failure of an attempt to prevent failures through specific data collection and processing solutions (e.g., templates), and the like.

[1126] In example embodiments, methods and systems of back-calculation of data collected with a system for data collection in an industrial environment may be applied to concrete pouring equipment in a construction site application. Concrete pouring equipment may comprise several active components including mixers that may include water and aggregate supply systems, mixing control systems, mixing motors, directional controllers, concrete sensors and the like, concrete pumps, delivery systems, flow control as well as on/off controls, and the like. Back-calculation of failure or other conditions of active or passive components of a concrete pouring equipment may benefit from visualization of the equipment, its components, sensors, and other points where data is collected (e.g., controllers and the like). Visualizing data/conditions collected from sensors associated with concrete pumps and the like when performing back-calculation of a flow rate failure condition may inform the user of a conditions of the pump that may contribute to the flow rate failure. Flow rate may decrease contemporaneously with an increase in temperature of the pump. This may be visualized by, for example, presenting the flow rate sensor data and the pump

temperature sensor data in the user interface. This correlation may be noted by an expert system or by a user observing the visualization and corrective action may be taken.

[1127] In example embodiments, methods and systems of back-calculation of data collected with a system for data collection in an industrial environment may be applied to digging and extraction systems in a mining application. Digging and extraction systems may comprise several active sub-systems including cutting heads, pneumatic drills, jack hammers, excavators, transport systems, and the like. Back-calculation of failure or other conditions of active or passive components of digging and extraction systems may benefit from visualization of the equipment, its components, sensors, and other points where data is collected (e.g., controllers and the like). Visualizing data/conditions collected from sensors associated with pneumatic drills and the like when performing back-calculation of a pneumatic line failure condition may inform the user of a conditions of the drill that may contribute to the line failure. Line pressure may increase contemporaneously with a change of a condition of the drill. This may be visualized by, for example, presenting the line pressure sensor data and data from sensors associated with the drill in the user interface. This correlation may be noted by an expert system or by a user observing the visualization and corrective action may be taken.

[1128] In example embodiments, methods and systems of back-calculation of data collected with a system for data collection in an industrial environment may be applied to cooling towers in an oil and gas production environment. Cooling towers may comprise several active components including feedwater systems, pumps, valves, temperature-controlled operation, storage systems, mixing systems, and the like. Back-calculation of failure or other conditions of active or passive components of cooling towers may benefit from visualization of the equipment, its components, sensors and other points where data is collected (e.g., controllers and the like). Visualizing data/conditions collected from sensors associated with the cooling towers and the like when performing back-calculation of a circulation pump failure condition may inform the user of a conditions of the cooling towers that may contribute to the pump failure. Temperature of the feedwater may increase contemporaneously with a decrease in output of the circulation pump. This may be visualized by, for example, presenting the feed water temperature sensor data and the pump output rate sensor data in the user interface. This correlation may be noted by an expert system or by a user observing the visualization and corrective action may be taken.

[1129] In example embodiments, methods and systems of back-calculation of data collected with a system for data collection in an industrial environment may be applied to circulation water systems in a power generation application. Circulation water systems may comprise several active components including, pumps, storage systems, water coolers, and the like. Back-calculation of failure or other conditions of active or passive components of circulation water systems may benefit from visualization of the equipment, its components, sensors and other points where data is collected (e.g., controllers and the like). Visualizing data/conditions collected from sensors associated with water coolers and the like when performing back-calculation of a circulation water temperature failure condition may inform the user of a conditions of the cooler that may contribute to the temperature condition failure. Circulation temperature may increase contemporaneously

with an increase of core water cooler temperature. This may be visualized by, for example, presenting the circulation water temperature sensor data and the water cooler temperature sensor data in the user interface. This correlation may be noted by an expert system or by a user observing the visualization and corrective action may be taken.

[1130] Referring to Figure 112 a graphical approach 11300 for back-calculation is depicted. Components of an industrial environment may be depicted in a map of the environment 11302. Components that may have a history of failure (with this installation or others) may be highlighted. In response to a selection of one of these components (such as by a user making the selection), related components and sensors for the selected part and related components may be highlighted, including signal routing paths for the data from their relevant sensors to a data collector. Additional highlighting may be added to sensors from which unacceptable data has been collected, thereby indicating potential root causes of a failure of the selected part. The relationships among the parts may be based at least in part on machine configuration metadata. The relationship between specific sensors and the failure condition may be based at least in part on a data collection template associated with the part and/or associated with the failure condition.

[1131] In example embodiments, a system for data collection in an industrial environment may route data from a plurality of sensors in the industrial environment to wearable haptic stimulators that present the data from the sensors as human detectable stimuli including at least one of tactile, vibration, heat, sound, and force. In example embodiments, the haptic stimulus represents an effect on the machine resulting from the sensed data. In example embodiments, a bending effect may be presented as bending a finger of a haptic glove. In example embodiments, a vibrating effect may be presented as vibrating a haptic arm band. In example embodiments, a heating effect may be presented as an increase in temperature of a haptic wrist band. In example embodiments, an electrical effect (e.g., over voltage, current, and others) may be presented as a change in sound of a phatic audio system.

[1132] In example embodiments, an industrial machine operator haptic user interface may be adapted to provide haptic stimuli to the operator that is responsive to the operator's control of the machine, wherein the stimuli indicate an impact on the machine as a result of the operator's control and interaction with objects in the environment as a result thereof. In example embodiments, sensed conditions of the machine that exceed an acceptable range may be presented to the operator through the haptic user interface. In example embodiments, the sensed conditions of the machine that are within an acceptable range may not be presented to the operator through the haptic user interface. In example embodiments, the sensed conditions of the machine that are within an acceptable range may presented as natural language representations of confirmation of the operator control. In example embodiments, at least a portion of the haptic user interface is worn by the operator. In example embodiments, a wearable haptic user interface device may include force exerting devices along the outer legs of a device operator's uniform. When a vehicle that the operator is controlling approaches an obstacle along a lateral side of the vehicle, an inflatable bellows may be inflated, exerting pressure against the leg of the operator closest to the side of the vehicle approaching the obstacle. The bellows may continue to be inflated, thereby exerting

additional pressure on the operator's leg that is consistent with the proximity of the obstacle. The pressure may be pulsed when contact with the obstacle is imminent. In another example, an arm band of an operator may vibrate in coordination with vibration being experienced by a portion of the vehicle that the operator is controlling. These are merely examples and not intended to be limiting or restrictive of the ways in which a wearable haptic feedback user device may be controlled to indicate conditions that are sensed by a system for data collection in an industrial environment.

[1133] In example embodiments, a haptic user interface safety system worn by a user in an industrial environment may be adapted to indicate proximity to the user of equipment in the environment by stimulating a portion of the user with at least one of pressure, heat, impact, electrical stimuli and the like, the portion of the user being stimulated may be closest to the equipment. In example embodiments, at least one of the type, strength, duration, and frequency of the stimuli is indicative of a risk of injury to the user.

[1134] In example embodiments, a wearable haptic user interface device, that may be worn by a user in an industrial environment, may broadcast its location and related information upon detection of an alert condition in the industrial environment. The alert condition may be proximal to the user wearing the device, or not proximal but related to the user wearing the device. The user's haptic device may broadcast the user's location to facilitate rapid access to the user or by the user to the emergency location when a user may be an emergency responder or there is detection of a situation requiring an emergency responder. In example embodiments, an alert condition may be determined from monitoring industrial machine sensors and may be presented to the user as haptic stimuli, with the severity of the alert corresponding to a degree of stimuli. In example embodiments, the degree of stimuli may be based on the severity of the alert, the corresponding stimuli may continue, be repeated, or escalate, optionally including activating multiple stimuli concurrently, send alerts to additional haptic users, and the like until an acceptable response is detected, e.g., through the haptic UI. The wearable haptic user device may be adapted to communicate with other haptic user devices to facilitate detecting the acceptable response.

[1135] In example embodiments, a wearable haptic user interface for use in an industrial environment may include gloves, rings, wrist bands, watches, arm bands, head gear, belts, necklaces, shirts (e.g., uniform shirt), foot wear, pants, ear protectors, safety glasses, vests, overalls, coveralls, and any other article of clothing or accessory that can be adapted to provide haptic stimuli.

[1136] In example embodiments, wearable haptic device stimuli may be correlated to a sensor in an industrial environment. Non-limiting examples include a vibration of a wearable haptic device in response to vibration detected in an industrial environment; increasing or decreasing the temperature of a wearable haptic device in response to a detected temperature in an industrial environment; producing sound that changes in pitch responsively to changes in a sensed electrical signal, and the like. In example embodiments, a severity of wearable haptic device stimuli may correlate to an aspect of a sensed condition in the industrial environment. Non-limiting examples include moderate or short-term vibration for a low degree of sensed vibration; strong or long-term

vibration stimulation for an increase in sensed vibration; aggressive, pulsed, and/or multi-mode stimulation for a high amount of sensed vibration. Wearable haptic device stimuli may also include lighting (e.g., flashing, color changes, and the like), sound, odor, tactile output, motion of the haptic device (e.g., inflating/deflating a balloon, extension/retraction of an articulated segment, and the like), force/impact, and the like.

[1137] In example embodiments, a system for data collection in an industrial environment may interface with wearable haptic feedback user devices to relay data collected from fuel handling systems in a power generation application to the user via haptic stimulation. Fuel handling for power generation may include solid fuels, such as woodchips, stumps, forest residue, sticks, energy willow, peat, pellets, bark, straw, agro biomass, coal, and solid recovery fuel. Handling systems may include receiving stations that may also sample the fuel, preparation stations that may crush or chip wood-based fuel or shred waste-based fuel. Fuel handling systems may include storage and conveying systems, feed and ash removal systems and the like. Wearable haptic user interface devices may be used with fuel handling systems by providing an operator feedback on conditions in the handling environment that the user is otherwise isolated from. Sensors may detect operational aspects of a solid fuel feed screw system. Conditions like screw rotational rate, weight of the fuel, type of fuel, and the like may be converted into haptic stimuli to a user while allowing the user to use his hands and provide his attention to operate the fuel feed system.

[1138] In example embodiments, a system for data collection in an industrial environment may interface with wearable haptic feedback user devices to relay data collected from suspension systems of a truck and/or vehicle application to the user via haptic stimulation. Haptic simulation may be correlated with conditions being sensed by the vehicle suspension system. In example embodiments, road roughness may be detected and converted into vibration-like stimuli of a wearable haptic arm band. In example embodiments, suspension forces (contraction and rebound) may be converted into stimuli that present a scaled down version of the forces to the user through a wearable haptic vest.

[1139] In example embodiments, a system for data collection in an industrial environment may interface with wearable haptic feedback user devices to relay data collected from hydroponic systems in an agriculture application to the user via haptic stimulation. In example embodiments, sensors in hydroponic systems, such as temperature, humidity, water level, plant size, carbon dioxide / oxygen levels, and the like may be converted to wearable device haptic stimuli. As an operator wearing haptic feedback clothing walks through a hydroponic agriculture facility, sensors proximal to the operator may signal to the haptic feedback clothing relevant information, such as temperature or a measure of actual temperature versus desired temperature that the haptic clothing may convert into haptic stimuli. In an example, a wrist band may include a thermal stimulator that can change temperature quickly to track temperature data or a derivative thereof from sensors in the agriculture environment. As a user walks through the facility, the haptic feedback wristband may change temperature to indicate a degree to which proximal temperatures are complying with expected temperatures.

[1140] In example embodiments, a system for data collection in an industrial environment may interface with wearable haptic feedback user devices to relay data collected from robotic positioning systems in an automated production line application to the user via haptic stimulation. Haptic feedback may include receiving a positioning system indicator of accuracy and converting it to an audible signal when the accuracy is acceptable, and another type of stimuli when the accuracy is not acceptable.

[1141] Referring to Figure 113, a wearable haptic user interface device for providing haptic stimuli to a user that is responsive to data collected in an industrial environment by a system adapted to collect data in the industrial environment is depicted. A system for data collection 11402 in an industrial environment 11400 may include a plurality of sensors. Data from those sensors may be collected and analyzed by a computing system. A result of the analysis may be communicated wirelessly to one or more wearable haptic feedback stimulators 11404 worn by a user associated with the industrial environment. The wearable haptic feedback stimulators may interpret the result, convert it into a form of stimuli based on a haptic stimuli-to-sensed condition mapping, and produce the stimuli.

[1142] In example embodiments, a system for data collection in an industrial environment may facilitate presenting a graphical element indicative of industrial machine sensed data on an augmented reality (AR) display. The graphical element may be adapted to represent a position of the sensed data on a scale of acceptable values of the sensed data. The graphical element may be positioned proximal to a sensor detected in the field of view being augmented that captured the sensed data in the AR display. The graphical element may be a color and the scale may be a color scale ranging from cool colors (e.g., greens, blues) to hot colors (e.g., yellow, red) and the like. Cool colors may represent data values closer to the midpoint of the acceptable range and the hot colors representing data values close to or outside of a maximum or minimum value of the range.

[1143] In example embodiments, a system for data collection in an industrial environment may present, in an AR display, data being collected from a plurality of sensors in the industrial environment as one of a plurality graphical effects (e.g., colors in a range of colors) that correlate the data being collected from each sensor to a scale of values within an acceptable range compared to values outside of the acceptable range. In example embodiments, the plurality of graphical effects may overlay a view of the industrial environment and placement of the plurality of graphical effects may correspond to locations in the view of the environment at which a sensor is located that is producing the corresponding sensor data. In example embodiments, a first set of graphical effects (e.g., hot colors) represent components for which multiple sensors indicate values outside acceptable ranges.

[1144] In example embodiments, a system for data collection in an industrial environment may facilitate presenting, in an AR display information being collected by sensors in the industrial environment as a heat map overlaying a visualization of the environment so that regions of the environment with sensor data suggestive of a greater potential of failure are overlaid with a graphic effect that is different than regions of the environment with sensor data suggestive of a lesser potential of failure. In example embodiments, the heat map is based on data currently being sensed.

In example embodiments, the heat map is based on data from prior failures. In example embodiments, the heat map is based on changes in data from an earlier period, such as data that suggest an increased likelihood of machine failure. In example embodiments, the heat map is based on a preventive maintenance plan and a record of preventive maintenance in the industrial environment.

[1145] In example embodiments, a system for data collection in an industrial environment may facilitate presenting information being collected by sensors in the industrial environment as a heat map overlaying a view of the environment, such as a live view as may be presented in an AR display. Such a system may include presenting an overlay that facilitates a call to action, wherein the overlay is associated with a region of the heat map. The overlay may comprise a visual effect of a part or subsystem of the environment on which the action is to be performed. In example embodiments, the action to be performed is maintenance related and may be part-specific.

[1146] In example embodiments, a system for data collection in an industrial environment may facilitate updating, in an AR view of a portion of the environment, a heat map of aspects of the industrial environment based on a change to operating instructions for at least one aspect of a machine in the industrial environment. The heat map may represent compliance with operational limits for portions of machines in the industrial environment. In example embodiments, the heat map may represent a likelihood of component failure as a result of the change to operation instructions.

[1147] In example embodiments, a system for data collection in an industrial environment may facilitate presenting, as a heat map in an AR view of a portion of the environment, a degree or measure of coverage of sensors in the industrial environment for a data collection template that identifies select sensors in the industrial environment for a data collection activity.

[1148] In example embodiments, a system for data collection in an industrial environment may facilitate displaying a heat map overlaying a view, such as a live view, of an industrial environment of failure-related data for various portions of the environment. The failure-related data may comprise a difference between an actual failure rate of the various portions and another failure rate. Another failure rate may be a rate of failure of comparable portions elsewhere in the environment, and/or average failure rate of comparable portions across a plurality of environments, such as an industry average, manufacturer failure rate estimate, and the like.

[1149] In example embodiments, a system for data collection in an industrial environment may facilitate displaying a heat map related to data collected from robotic arms and hands for production line robotic handling in an augmented reality view of a portion of the environment. A heat map related to data collected from robotic arms and hands may represent data from sensors disposed in – for example, the fingers of a robotic hand. Sensor may collect data, such as applied pressure when pinching an object, resistance (e.g., responsive to a robotic touch) of an object, multi-axis forces presented to the finger as it performs an operation, such as holding a tool and the like, temperature of the object, total movement of the finger from initial point of contact until a resistance threshold is met, and other hand position / use conditions. Heat maps of this data may be presented in an augmented reality view of a robotic production environment so that a user may

make a visual assessment of, for example, how the relative positioning of the robotic fingers impacts the object being handled.

[1150] In example embodiments, a system for data collection in an industrial environment may facilitate displaying a heat map related to data collected from linear bearings for production line robotic handling in an augmented reality view of a portion of the environment. Linear bearings, as with most bearings, may not be visible while in use. However, assessing their operation may benefit from representing data from sensors that capture information about the bearings while in use in an augmented reality display. In example embodiments, sensors may be placed to detect forces being placed on portions of the bearings by the rotating member or elements that the bearings support. These forces may be presented as heat maps that correspond to relative forces, on a visualization of the bearings in an augmented reality view of a robot handling machine that uses linear bearings.

[1151] In example embodiments, a system for data collection in an industrial environment may facilitate displaying a heat map related to data collected from boring machinery for mining in an augmented reality view of a portion of the environment. Boring machinery, and in particular multi-tip circular boring heads may experience a range of rock formations at the same time. Sensors may be placed proximal to each boring tip that may detect forces experienced by the tips. The data may be collected by a system adapted to collect data in an industrial environment and provided to an augmented reality system that may display the data as heat maps or the like in a view of the boring machine.

[1152] Referring to Figure 114, an augmented reality display of heat maps based on data collected in an industrial environment by a system adapted to collect data in the environment is depicted. An augmented reality view of an industrial environment 11500 may include heat maps 11502 that depict data received from or derived from data received from sensors 11504 in the industrial environment. Sensor data may be captured and processed by a system adapted for data collection and analysis in an industrial environment. The data may be converted into a form that is suitable for use in an augmented reality system for displaying heat maps. The heat maps 11502 may be aligned in the augmented reality view with a sensor from which the underlying data was sourced.

[1153] In example embodiments, a system for data collection and visualization thereof in an industrial environment may include an augmented reality and/or virtual reality (AR/VR) display in which data values output by sensors disposed in a field of view in the AR/VR display are displayed with visual attributes that indicate a degree of compliance of the data to an acceptable range or values for the sensed data. In example embodiments, the visual attributes may provide near real-time portrayal of trends of the sensed data and/or of derivatives thereof. In example embodiments, the visual attributes may be the actual data being captured, or the derived data, such as a trend of the data and the like.

[1154] In example embodiments, a system for data collection and visualization thereof in an industrial environment may include an AR/VR display in which trends of data values output by sensors disposed in a field of view in the AR/VR are displayed with visual attributes that indicate a degree of severity of the trend. In example embodiments, other data or analysis that could be

displayed may include: data from sensors that exceed an acceptable range, data from sensors that are part of a smart band selected by the user, data from sensors that are monitored for triggering a smart band collection action, data from sensors that sense an aspect of the environment that meets preventive maintenance criteria, such as a PM action is upcoming soon, a PM action was recently performed or is overdue for PM. Other data for such AR/VR visualization may include data from sensors for which an acceptable range has recently been changed, expanded, narrowed and the like. Other data for such AR/VR visualization that may be particularly useful for an operator of an industrial machine (digging, drilling, and the like) may include analysis of data from sensors, such as for example impact on an operating element (torque, force, strain, and the like).

[1155] In example embodiments, a system for data collection and visualization thereof in an industrial environment that may include presentation of visual attributes that represent collected data in an AR/VR environment may do so for pumps in a mining application. Mining application pumps may provide water and remove liquefied waste from a mining site. Pump performance may be monitored by sensors detecting pump motors, regulators, flow meters, and the like. Pump performance monitoring data may be collected and presented as a set of visual attributes in an augmented reality display. In an example, pump motor power consumption, efficiency, and the like may be displayed proximal to a pump viewed through an augmented reality display.

[1156] In example embodiments, a system for data collection and visualization thereof in an industrial environment that may include presentation of visual attributes that represent collected data in an AR/VR environment may do so for energy storage in a power generation application. Power generation energy storage may be monitored with sensors that capture data related to storage and use of stored energy. Information such as utilization of individual energy storage cells, energy storage rate (e.g., battery charging and the like), stored energy consumption rate (e.g., KWH being supplied by an energy storage system), storage cell status, and the like may be captured and converted into augmented reality viewable attributes that may be presented in an augmented reality view of an energy storage system.

[1157] In example embodiments, a system for data collection and visualization thereof in an industrial environment that may include presentation of visual attributes that represent collected data in an AR/VR environment may do so for feed water systems in a power generation application. Sensors may be disposed in an industrial environment, such as power generation for collecting data about feed water systems. Data from those sensors may be captured and processed by the system for data collection. Results of this processing may include trends of the data, such as feed water cooling rates, flow rates, pressure and the like. These trends may be presented on an augmented reality view of a feed water system by applying a map of sensors with physical elements visible in the view and then retrieving data from the mapped sensors. The retrieved data (and derivatives thereof) may be presented in the augmented reality view of the feed water system.

[1158] Referring to Figure 115, an augmented reality display 11600 comprising real time data 11602 overlaying a view of an industrial environment is depicted. Sensors 11604 in the environment may be recognized by the augmented reality system, such as by first detecting an industrial machine, system, or part thereof with which the sensors are associated. Data from the

sensors 11604 may be retrieved from a data repository, processed into trends, and presented in the augmented reality view 11600 proximal to the sensors from which the data originates.

[1159] Figure 158 shows a system for data collection in an industrial environment having a self-sufficient data acquisition box for capturing and analyzing data in an industrial environment including sensor inputs 11700, 11702, 11704, 11706 that connect to a data circuit 11708 for analyzing the sensor inputs, a network communication interface 11712, a network control circuit 11710 for sending and receiving information related to the sensor inputs to an external system and a data filter circuit configured to dynamically adjust what portion of the information is sent based on instructions received over the network communication interface. A variety of sensor inputs X connect to the data circuit Y. The data circuit intercommunicates with a network control circuit, which is connected to one or more network interfaces. These interfaces may include wired interfaces or wireless interfaces, communicating via a star, multi-hop, peer-to-peer, hub-and-spoke, mesh, ring, hierarchical, daisy-chained, broadcast, or other networking protocol. These interfaces may be multi-pair as in Ethernet, or single-wire networking protocol such as I2C. The networking protocol may interface one or more of a variety of variants of Ethernet and other protocols for real-time communication in an industrial network, including Modbus® over TCP, Industrial Ethernet, Ethernet Powerlink, Ethernet/IP, EtherCAT, Sercos®, Profinet™, CAN bus, serial protocols, near-field protocols, as well as home automation protocols such as ZIGBEE®, Z-WAVE™, or wireless WWAN or WLAN protocols such as LTE™, Wi-Fi, BLUETOOTH™, or others. The sensor inputs can be permanently or removably connected to the thing they are measuring, or may be integrated in a standalone data acquisition box. The entire system may be integrated into the apparatus that is being measured, such as a vehicle (e.g., a car, a truck, a commercial vehicle, a tractor, a construction vehicle or other type of vehicle), a component or item of equipment (e.g., a compressor, agitator, motor, fan, turbine, generator, conveyor, lift, robotic assembly, or any other item as described throughout this disclosure), an infrastructure element (such as a foundation, a housing, a wall, a floor, a ceiling, a roof, a doorway, a ramp, a stairway, or the like) or other feature or aspect of an industrial environment. The entire system may be integrated into a stationary industrial system such as a production assembly, static components of an assembly line subject to wear and stress (such as rail guides), or motive elements such as robotics, linear actuators, gearboxes, and vibrators.

[1160] Disclosed herein are methods and systems for data collection in an industrial environment featuring self-organization functionality. Such data collection systems and methods may facilitate intelligent, situational, context-aware collection, summarization, storage, processing, transmitting, and/or organization of data, such as by one or more data collectors (such as any of the wide range of data collector embodiments described throughout this disclosure), a central headquarters or computing system, and the like. The described self-organization functionality of data collection in an industrial environment may improve various parameters of such data collection, as well as parameters of the processes, applications, and products that depend on data collection, such as data quality parameters, consistency parameters, efficiency parameters, comprehensiveness parameters, reliability parameters, effectiveness parameters, storage utilization parameters, yield parameters

(including financial yield, output yield, and reduction of adverse events), energy consumption parameters, bandwidth utilization parameters, input/output speed parameters, redundancy parameters, security parameters, safety parameters, interference parameters, signal-to-noise parameters, statistical relevancy parameters, and others. The self-organization functionality may optimize across one or more such parameters, such as based on a weighting of the value of the parameters; for example, a swarm of data collectors may be managed (or manage itself) to provide a given level of redundancy for critical data, while not exceeding a specified level of energy usage, e.g., per data collector or a group of data collectors or the entire swarm of data collectors. This may include using a variety of optimization techniques described throughout this disclosure and the documents incorporated herein by reference.

[1161] In example embodiments, such methods and systems for data collection in an industrial environment can include one or more data collectors, e.g., arranged in a cooperative group or "swarm" of data collectors, that collect and organize data in conjunction with a data pool in communication with a computing system, as well as supporting technology components, services, processes, modules, applications and interfaces, for managing the data collection (collectively referred to in some cases as a data collection system 12004). Examples of such components include, but are not limited to, a model-based expert system, a rule-based expert system, an expert system using artificial intelligence (such as a machine learning system, which may include a neural net expert system, a self-organizing map system, a human-supervised machine learning system, a state determination system, a classification system, or other artificial intelligence system), or various hybrids or combinations of any of the above. References to a self-organizing method or system should be understood to encompass utilization of any one of the foregoing or suitable combinations, except where context indicates otherwise.

[1162] The data collection systems and methods of the present disclosure can be utilized with various types of data, including but not limited to vibration data, noise data and other sensor data of the types described throughout this disclosure. Such data collection can be utilized for event detection, state detection, and the like, and such event detection, state detection, and the like can be utilized to self-organize the data collection systems and methods, as further discussed herein. The self-organization functionality may include managing data collector(s), both individually or in groups, where such functionality is directed at supporting an identified application, process, or workflow, such as confirming progress toward or /alignment with one or more objectives, goals, rules, policies, or guidelines. The self-organization functionality may also involve managing a different goal/guideline, or directing data collectors targeted to determining an unknown variable based on collection of other data (such as based on a model of the behavior of a system that involves the variable), selecting preferred sensor inputs among available inputs (including specifying combinations, fusions, or multiplexing of inputs), and/or specifying a specific data collector among available data collectors.

[1163] A data collector may include any number of items, such as sensors, input channels, data locations, data streams, data protocols, data extraction techniques, data transformation techniques, data loading techniques, data types, frequency of sampling, placement of sensors, static data points,

metadata, fusion of data, multiplexing of data, self-organizing techniques, and the like as described herein. Data collector settings may describe the configuration and makeup of the data collector, such as by specifying the parameters that define the data collector. For example, data collector settings may include one or more frequencies to measure. Frequency data may further include at least one of a group of spectral peaks, a true-peak level, a crest factor derived from a time waveform, and an overall waveform derived from a vibration envelope, as well as other signal characteristics described throughout this disclosure. Data collectors may include sensors measuring or data regarding one or more wavelengths, one or more spectra, and/or one or more types of data from various sensors and metadata. Data collectors may include one or more sensors or types of sensors of a wide range of types, such as described throughout this disclosure and the documents incorporated by reference herein. Indeed, the sensors described herein may be used in any of the methods or systems described throughout this disclosure. For example, one sensor may be an accelerometer, such as one that measures voltage per G of acceleration (e.g., 100 mV/G, 500 mV/G, 1 V/G, 5 V/G, 10 V/G). In example embodiments, a data collector may alter the makeup of the subset of the plurality of sensors used in a data collector based on optimizing the responsiveness of the sensor, such as for example choosing an accelerometer better suited for measuring acceleration of a lower speed gear system or drill/boring device versus one better suited for measuring acceleration of a higher speed turbine in a power generation environment. Choosing may be done intelligently, such as for example with a proximity probe and multiple accelerometers disposed on a specific target (e.g., a gear system, drill, or turbine) where while at low speed one accelerometer is used for measuring in the data collector and another is used at high speeds. Accelerometers come in various types, such as piezo-electric crystal, low frequency (e.g., 10 V/G), high speed compressors (10 MV/G), MEMS, and the like. In another example, one sensor may be a proximity probe which can be used for sleeve or tilt-pad bearings (e.g., oil bath), or a velocity probe. In yet another example, one sensor may be a solid state relay (SSR) that is structured to automatically interface with another routed data collector (such as a mobile or portable data collector) to obtain or deliver data. In another example, a data collector may be routed to alter the makeup of the plurality of available sensors, such as by bringing an appropriate accelerometer to a point of sensing, such as on or near a component of a machine. In still another example, one sensor may be a triax probe (e.g., a 100 MV/G triax probe), that in embodiments is used for portable data collection. In example embodiments, of a triax probe, a vertical element on one axis of the probe may have a high frequency response while the ones mounted horizontally may influence limit the frequency response of the whole triax. In another example, one sensor may be a temperature sensor and may include a probe with a temperature sensor built inside, such as to obtain a bearing temperature. In still additional examples, sensors may be ultrasonic, microphone, touch, capacitive, vibration, acoustic, pressure, strain gauges, thermographic (e.g., camera), imaging (e.g., camera, laser, IR, structured light), a field detector, an EMF meter to measure an AC electromagnetic field, a gaussmeter, a motion detector, a chemical detector, a gas detector, a CBRNE detector, a vibration transducer, a magnetometer, positional, location-based, a velocity sensor, a displacement sensor, a tachometer, a flow sensor, a level sensor, a proximity sensor, a pH

sensor, a hygrometer / moisture sensor, a densitometric sensor, an anemometer, a viscometer, or any analog industrial sensor and/or digital industrial sensor. In a further example, sensors may be directed at detecting or measuring ambient noise, such as a sound sensor or microphone, an ultrasound sensor, an acoustic wave sensor, and an optical vibration sensor (e.g., using a camera to see oscillations that produce noise). In still another example, one sensor may be a motion detector.

[1164] Data collectors may be of or may be configured to encompass one or more frequencies, wavelengths or spectra for particular sensors, for particular groups of sensors, or for combined signals from multiple sensors (such as involving multiplexing or sensor fusion). Data collectors may be of or may be configured to encompass one or more sensors or sensor data (including groups of sensors and combined signals) from one or more pieces of equipment/components, areas of an installation, disparate but interconnected areas of an installation (e.g., a machine assembly line and a boiler room used to power the line), or locations (e.g., a building in one geographic location and a building in a separate, different geographic location). Data collector settings, configurations, instructions, or specifications (collectively referred to herein using any one of those terms) may include where to place a sensor, how frequently to sample a data point or points, the granularity at which a sample is taken (e.g., a number of sampling points per fraction of a second), which sensor of a set of redundant sensors to sample, an average sampling protocol for redundant sensors, and any other aspect that would affect data acquisition.

[1165] Within the data collection system 12004, the self-organization functionality can be implemented by a neural net, a model-based system, a rule-based system, a machine learning system, and/or a hybrid of any of those systems. Further, the self-organizing functionality may be performed in whole or in part by individual data collectors, a collection or group of data collectors, a network-based computing system, a local computing system comprising one or more computing devices, a remote computing system comprising one or more computing devices, and a combination of one or more of these components. The self-organization functionality may be optimized for a particular goal or outcome, such as predicting and managing performance, health, or other characteristics of a piece of equipment, a component, or a system of equipment or components. Based on continuous or periodic analysis of sensor data, as patterns/trends are identified, or outliers appear, or a group of sensor readings begin to change, etc., the self-organization functionality may modify the collection of data intelligently, as described herein. This may occur by triggering a rule that reflects a model or understanding of system behavior (e.g., recognizing a shift in operating mode that calls for different sensors as velocity of a shaft increases) or it may occur under control of a neural net (either in combination with a rule-based approach or on its own), where inputs are provided such that the neural net over time learns to select appropriate collection modes based on feedback as to successful outcomes (e.g., successful classification of the state of a system, successful prediction, successful operation relative to a metric). For example only, when an assembly line is reconfigured for a new product or a new assembly line is installed in a manufacturing facility, data from the current data collector(s) may not accurately predict the state or metric of operation of the system, thus, the self-organization functionality may begin to

iterate to determine if a new data collector, type of sensed data, format of sensed data, etc. is better at predicting a state or metric. Based on offset system data, such as from a library or other data structure, certain sensors, frequency bands or other data collectors may be used in the system initially and data may be collected to assess performance. As the self-organization functionality iterates, other sensors/frequency bands may be accessed to determine their relative weight in identifying performance metrics. Over time, a new frequency band may be identified (or a new collection of sensors, a new set of configurations for sensors, or the like) as a better or more suitable gauge of performance in the system and the self-organization functionality may modify its data collector(s) based on this iteration. For example only, perhaps an older boring tool in an energy extraction environment dampens one or more vibration frequencies while a different frequency is of higher amplitude and present during optimal performance than what was seen in the present system. In this example, the self-organization functionality may alter the data collectors from what was originally proposed, e.g., by the data collection system, to capture the higher amplitude frequency that is present in the current system.

[1166] The self-organization functionality, in embodiments involving a neural net or other machine learning system, may be seeded and may iterate, e.g., based on feedback and operation parameters, such as described herein. Certain feedback may include utilization measures, efficiency measures (e.g., power or energy utilization, use of storage, use of bandwidth, use of input/output use of perishable materials, use of fuel, and/or financial efficiency, financial such as reduction of costs), measures of success in prediction or anticipation of states (e.g., avoidance and mitigation of faults), productivity measures (e.g., workflow), yield measures, and profit measures. Certain parameters may include storage parameters (e.g., data storage, fuel storage, storage of inventory), network parameters (e.g., network bandwidth, input/output speeds, network utilization, network cost, network speed, network availability), transmission parameters (e.g., quality of transmission of data, speed of transmission of data, error rates in transmission, cost of transmission), security parameters (e.g., number and/or type of exposure events, vulnerability to attack, data loss, data breach, access parameters), location and positioning parameters (e.g., location of data collectors, location of workers, location of machines and equipment, location of inventory units, location of parts and materials, location of network access points, location of ingress and egress points, location of landing positions, location of sensor sets, location of network infrastructure, location of power sources) , input selection parameters, data combination parameters (e.g., for multiplexing, extraction, transformation, loading), power parameters (e.g., of individual data collectors, groups of data collectors, or all potentially available data collectors), states (e.g., operational modes, availability states, environmental states, fault modes, health states, maintenance modes, anticipated states), events, and equipment specifications. With respect to states, operating modes may include, mobility modes (direction, speed, acceleration, and the like), type of mobility modes (e.g., rolling, flying, sliding, levitation, hovering, floating), performance modes (e.g., gears, rotational speeds, heat levels, assembly line speeds, voltage levels, frequency levels), output modes, fuel conversion modes, resource consumption modes, and financial performance modes (e.g., yield, profitability). Availability states may refer to anticipating

conditions that could cause one or more machines to go offline or require backup. Environmental states may refer to ambient temperature, ambient humidity/moisture, ambient pressure, ambient wind/fluid flow, presence of pollution or contaminants, presence of interfering elements (e.g., electrical noise, vibration), power availability, and power quality, among other parameters. Anticipated states may include achieving or not achieving a desired goal, such as a specified/threshold output production rate, a specified/threshold generation rate, an operational efficiency/failure rate, a financial efficiency/profit goal, a power efficiency/resource utilization, an avoidance of a fault condition (e.g., overheating, slow performance, excessive speed, excessive motion, excessive vibration/oscillation, excessive acceleration, expansion/contraction, electrical failure, running out of stored power/fuel, overpressure, excessive radiation/melt down, fire, freezing, failure of fluid flow (e.g., stuck valves, frozen fluids), mechanical failures (e.g., broken component, worn component, faulty coupling, misalignment, asymmetries/deflection, damaged component (e.g., deflection, strain, stress, cracking), imbalances, collisions, jammed elements, and lost or slipping chain or belt), avoidance of a dangerous condition or catastrophic failure, and availability (online status)).

[1167] The self-organization functionality may comprise or be seeded with a model that predicts an outcome or state given a set of data, which may comprise inputs from sensors, such as via a data collector, as well as other data, such as from system components, from external systems and from external data sources. For example, the model may be an operating model for an industrial environment, machine, or workflow. In another example, the model may be for anticipating states, for predicting fault, for optimizing maintenance, for optimizing data transport (such as for optimizing network coding, network-condition-sensitive routing), for optimizing data marketplaces, and the like.

[1168] The self-organization functionality may result in any number of downstream actions based on analysis of data from the data collector(s). In example embodiments, the self-organization functionality may determine that the system should either keep or modify operational parameters, equipment or a weighting of a neural net model given a desired goal, such as a specified/threshold output production rate, specified/threshold generation rate, an operational efficiency/failure rate, a financial efficiency/profit goal, a power efficiency/resource utilization, an avoidance of a fault condition, an avoidance of a dangerous condition or catastrophic failure, and the like. In example embodiments, the adjustments may be based on determining context of an industrial system, such as understanding a type of equipment, its purpose, its typical operating modes, the functional specifications for the equipment, the relationship of the equipment to other features of the environment (including any other systems that provide input to or take input from the equipment), the presence and role of operators (including humans and automated control systems), and ambient or environmental conditions. For example, in order to achieve a profit goal in a distribution environment (e.g., a power distribution environment), a generator or system of generators may need to operate at a certain efficiency level. The self-organization functionality may be seeded with a model for operation of the system of generators in a manner that results in a specified profit goal, such as indicating an on/off state for individual generator(s) in the power generation system based

on the time of day, current market sale price for the fuel consumed by the generators, current demand or anticipated future demand, and the like. As it acquires data and iterates, the model predicts whether the profit goal will be achieved given the current data, and determines whether the data or type of data being collected is appropriate, sufficient, etc. for the model. Based on the results of the iteration, a recommendation may be made (or a control instruction may be automatically provided) to gather different/additional data, organize the data differently, direct different data collectors to collect new data, etc. and/or to operate a subset of the generators at a higher output (but less efficient) rate, power on additional generators, maintain a current operational state, or the like. Further, as the system iterates, one or more additional sensors may be sampled in the model to determine if their addition to the self-organization functionality would improve predicting a state or otherwise assisting with the goals of the data collection efforts.

[1169] In example embodiments, a system for data collection in an industrial environment may include a plurality of input sensors, such as any of those described herein, communicatively coupled to a data collector having one or more processors. The data collection system may include a plurality of individual data collectors structured to operate together to determine at least one subset of the plurality of sensors from which to process output data. The data collection system may also include a machine learning circuit structured to receive output data from the at least one subset of the plurality of sensors and learn received output data patterns indicative of a state. In example embodiments, the data collection system may alter the at least one subset of the plurality of sensors, or an aspect thereof, based on one or more of the learned received output data patterns and the state. In certain example embodiments, the machine learning circuit is seeded with a model that enables it to learn data patterns. The model may be a physical model, an operational model, a system model, and the like. In other embodiments, the machine learning circuit is structured for deep learning wherein input data is fed to the circuit with no or minimal seeding and the machine learning data analysis circuit learns based on output feedback. For example, a metal tooling system in a manufacturing environment may operate to manufacture parts using machine tools such as lathes, milling machines, grinding machines, boring tools, and the like. Such machines may operate at various speeds and output rates, which may affect the longevity, efficiency, accuracy, etc. of the machine. The data collector may acquire various parameters to evaluate the environment of the machine tools, e.g., speed of operation, heat generation, vibration, and conformity with a part specification. The system can utilize such parameters and iterate towards a prediction of state, output rate, etc. based on such feedback. Further, the system may self-organize such that the data collector(s) collect additional/different data from which such predictions may be made.

[1170] There may be a balance of multiple goals/guidelines in the self-organization functionality of data collection system. For example, a repair and maintenance organization (RMO) may have operating parameters designed for maintenance of a machine in a manufacturing facility, while the owner of the facility may have particular operating parameters for the machine that are designed for meeting a production goal. These goals, in this example relating to a maintenance goal or a production output, may be tracked by a different data collectors or sensors. For example, maintenance of a machine may be tracked by sensors including a temperature sensor, a vibration

transducer, and a strain gauge while the production goal of a machine may be tracked by sensors including a speed sensor and a power consumption meter. The data collection system may (optionally using a neural net, machine learning system, deep learning system, or the like, which may occur under supervision by one or more supervisors (human or automated) intelligently manage data collectors aligned with different goals and assign weights, parameter modifications, or recommendations based on a factor, such as a bias towards one goal or a compromise to allow better alignment with all goals being tracked, for example. Compromises among the goals delivered to the data collection system may be based on one or more hierarchies or rules relating to the authority, role, criticality, or the like of the applicable goals. In example embodiments, compromises among goals may be optimized using machine learning, such as a neural net, deep learning system, or other artificial intelligence system as described throughout this disclosure. For example, in a power plant where a turbine is operating, the data collection system may manage multiple data collectors, such as one directed to detecting the operational status of the turbine, one directed at identifying a probability of hitting a production goal, and one directed at determining if the operation of the turbine is meeting a fuel efficiency goal. Each of these data collectors may be populated with different sensors or data from different sensors (e.g., a vibration transducer to indicate operational status, a flow meter to indicate production goal, and a fuel gauge to indicate a fuel efficiency) whose output data are indicative of an aspect of a particular goal. Where a single sensor or a set of sensors is helpful for more than one goal, overlapping data collectors (having some sensors in common and other sensors not in common) may take input from that sensor or set of sensors, as managed by the data collection system. If there are constraints on data collection (such as due to power limitations, storage limitations, bandwidth limitations, input/output processing capabilities, or the like), a rule may indicate that one goal (e.g., a fuel utilization goal or a pollution reduction goal that is mandated by law or regulation) takes precedence, such that the data collection for the data collectors associated with that goal are maintained as others are paused or shut down. Management of prioritization of goals may be hierarchical or may occur by machine learning. The data collection system may be seeded with models, or may not be seeded at all, in iterating towards a predicted state (e.g., meeting a goal) given the current data it has acquired. In this example, during operation of the turbine the plant owner may decide to bias the system towards fuel efficiency. All of the data collectors may still be monitored, but as the self-organization functionality iterates and predicts that the system will not collect or is not collecting data sufficient to determine whether the system is or is not meeting a particular goal, the data collection system may recommend or implement changes directed at collecting the appropriate data. Further, the plant owner may structure the system with a bias towards a particular goal such that the recommended changes to data collection parameters affecting such goal are made in favor of making other recommended changes.

[1171] In example embodiments, the data collection system may continue iterating in a deep-learning fashion to arrive at a distribution of data collectors, after being seeded with more than one data collection data type, that optimizes meeting more than one goal. For example, there may be multiple goals tracked for a refining environment, such as refining efficiency and economic

efficiency. Refining efficiency for the refining system may be expressed by comparing fuel put into the system, which can be obtained by knowing the amount of and quality of the fuel being used, and the amount of the refined product output from the system, which is calculated using the flow out of the system. Economic efficiency of the refining system may be expressed as the ratio between costs to run the system, including fuel, labor, materials and services, and the refined product output from the system for a period of time. Data used to track refining efficiency may include data from a flow meter, quality data point(s), and a thermometer, and data used to track economic efficiency may be a flow of product output from the system and costs data. These data may be used in the data collection system to predict states; however, the self-organization functionality of the system may iterate towards a data collection strategy that is optimized to predict states related to both thermal and economic efficiency. The new data collection schema may include data used previously in the individual data collectors but may also use new data from different sensors or data sources.

[1172] The iteration of the data collection system may be governed by rules, in some embodiments. For example, the data collection system may be structured to collect data for seeding at a pre-determined frequency. The data collection system may be structured to iterate at least a number of times, such as when a new component/equipment/fuel source is added, when a sensor goes off-line, or as standard practice. For example, when a sensor measuring the rotation of a boring tool in an offshore drilling operation goes off-line and the data collection system begins acquiring data from a new sensor or data collector measuring the same data points, the data collection system may be structured to iterate for a number of times before the state is utilized in or allowed to affect any downstream actions. The data collection system may be structured to train off-line or train in situ/online. The data collection system may be structured to include static and/or manually input data in its data collectors. For example, a data collection system associated with such a boring tool may be structured to iterate towards predicting a distance bored based on a duration of operation, wherein the data collector(s) include data regarding the speed of the boring tools, a distance sensor, a temperature sensor, and the like.

[1173] In example embodiments, the data collection system may be overruled. In example embodiments, the data collection system may revert to prior settings, such as in the event the self-organization functionality fails, such as if the collected data is insufficient or inappropriately collected, if uncertainty is too high in a model-based system, if the system is unable to resolve conflicting rules in rule-based system, or the system cannot converge on a solution in any of the foregoing. For example, sensor data on a power generation system used by the data collection system may indicate a non-operational state (such as a seized turbine), but output sensors and visual inspection, such as by a drone, may indicate normal operation. In this event, the data collection system may revert to an original data collection schema for seeding the self-organization functionality. In another example, one or more point sensors on a refrigeration system may indicate imminent failure in a compressor, but the data collector self-organized to collect data associated towards determining a performance metric did not identify the failure. In this event, the data

collector(s) will revert to an original setting or a version of the data collector setting that would have also identified the imminent failure of the compressor.

[1174] In example embodiments, the data collection system may change data collector settings in the event that a new component is added that makes the system closer to a different system. For example, a vacuum distillation unit is added to an oil and gas refinery to distill naphthalene, but the current data collector settings for the data collection system are derived from a refinery that distills kerosene. In this example, a data structure with data collector settings for various systems may be searched for a system that is more closely matched to the current system. When a new system is identified as more closely matched, such as one that also distill naphthalene, the new data collector settings (which sensors to use, where to direct them, how frequently to sample, what types of data and points are needed, etc. as described herein) are used to seed the data collection system to iterate towards predicting a state for the system. In example embodiments, the data collection system may change data collector settings in the event that a new set of data is available from a third party library. For example, a power generation plant may have optimized a specific turbine model to operate in a highly efficient way and deposited the data collector settings in a data structure. The data structure may be continuously scanned for new data collectors that better aid in monitoring power generation and thus, result in optimizing the operation of the turbine.

[1175] In example embodiments, the data collection system may utilize self-organization functionality to uncover unknown variables. For example, the data collection system may iterate to identify a missing variable to be used for further iterations. For example, an under-utilized tank in a legacy condensate/make-up water system of a power station may have an unknown capacity because it is inaccessible and no documentation exists on the tank. Various aspects of the tank may be measured by a swarm of data collectors to arrive at an estimated volume (e.g., flow into a downstream space, duration of a dye traced solution to work through the system), which can then be fed into the data collection system as a new variable.

[1176] In example embodiments, the data collection system node may be on a machine, on a data collector (or a group of them), in a network infrastructure (enterprise or other), or in the cloud. In example embodiments, there may be distributed neurons across nodes (e.g., machine, data collector, network, cloud).

[1177] In an aspect, and as illustrated in Figure 118, a data collection system 12004 can be arranged to collect data in an industrial environment 12000, e.g., from one or more targets 12002. In the illustrated embodiments, the data collection system 12004 includes a group or "swarm" 12006 of data collectors 12008, a network 12010, a computing system 12012, and a database or data pool 12014. Each of the data collectors 12008 can include one or more input sensors and be communicatively coupled to any and all of the other components of the data collection system 12004, as is partially illustrated by the connecting arrows between components.

[1178] The targets 12002 can be any form of machinery or component thereof in an industrial environment 12000. Examples of such industrial environments 12000 include but are not limited to factories, pipelines, construction sites, ocean oil rigs, ships, airplanes or other aircraft, mining environments, drilling environments, refineries, distribution environments, manufacturing

environments, energy source extraction environments, offshore exploration sites, underwater exploration sites, assembly lines, warehouses, power generation environments, and hazardous waste environments, each of which may include one or more targets 12002. Targets 12002 can take any form of item or location at which a sensor can obtain data. Examples of such targets 12002 include but are not limited to machines, pipelines, equipment, installations, tools, vehicles, turbines, speakers, lasers, automatons, computer equipment, industrial equipment, and switches.

[1179] The self-organization functionality of the data collection system 12004 can be performed at or by any of the components of the data collection system 12004. In example embodiments, a data collector 12008 or the swarm 12006 of data collectors 12008 can self-organize without assistance from other components and based on, e.g., the data sensed by its associated sensors and other knowledge. In example embodiments, the network 12010 can self-organize without assistance from other components and based on, e.g., the data sensed by the data collectors 12008 or other knowledge. Similarly, the computing system 12012 and/or the data pool 12014 without assistance from other components and based on, e.g., the data sensed by the data collectors 12008 or other knowledge. It should be appreciated that any combination or hybrid-type self-organization system can also be implemented.

[1180] For example only, the data collection system 12004 can perform or enable various methods or systems for data collection having self-organization functionality in an industrial environment 12000. These methods and systems can include analyzing a plurality of sensor inputs, e.g., received from or sensed by sensors at the data collector(s) 12008. The methods and systems can also include sampling the received data and self-organizing at least one of: (i) a storage operation of the data; (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1181] In aspects, the storage operation can include storing the data in a local database, e.g., of a data collector 12008, a computing system 12012, and/or a data pool 12014. The data can also be summarized over a given time period to reduce a size of the sensed data. The summarized data can be sent to one or more data acquisition boxes, to one or more data centers, and/or to other components of the system or other, separate systems. Summarizing the data over a given time period to reduce the size of the data, in some aspects, can include determining a speed at which data can be sent via a network (e.g., network 12010), wherein the size of the summarized data corresponds to the speed at which data can be sent continuously in real time via the network. In such aspects, or others, the summarized data can be continuously sent, e.g., to an external device via the network.

[1182] In various implementations, the methods and systems can include committing the summarized data to a local ledger, identifying one or more other accessible signal acquisition instruments on an accessible network, and/or synchronizing the summarized data at the local ledger with at least one of the other accessible signal acquisition instruments (e.g., data collectors 12008). In example embodiments, receiving a remote stream of sensor data from one or more other accessible signal acquisition instruments via a network can be included. An advertisement message

to a potential client indicating availability of at least one of the locally stored data, the summarized data, and the remote stream of sensor data can also or alternatively be sent.

[1183] The methods and systems can include identifying one or more other accessible signal acquisition instruments (e.g., data collectors 12008) on an accessible network (e.g., 12010), nominating at least one of the one or more other accessible signal acquisition instruments as a logical communication hub, and providing the logical communication hub with a list of available data and their associated sources. The list of available data and their associated sources can be provided to the logical communication hub utilizing a hybrid peer-to-peer communications protocol.

[1184] In some aspects, the storage operation can include storing the data in a local database and automatically organizing at least one parameter of the data pool utilizing machine learning. The organizing can be based at least in part on receiving information regarding at least one of an accuracy of classification and an accuracy of prediction of an external machine learning system that uses data from the data pool (e.g., data pool 12014).

[1185] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs and self-organizing at least one of: (i) a storage operation of the data; (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1186] The present disclosure describes a system for data collection in an industrial environment having self-organization functionality, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the industrial environment and for generating data associated with the plurality of sensor inputs, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1187] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs,

[1188] sampling data received from the sensor inputs; and self-organizing at least one of: (i) a storage operation of the data; (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the storage operation includes storing the data in a local database, and summarizing the data over a given time period to reduce a size of the data.

[1189] In example embodiments, the method further includes sending the summarized data to one or more data acquisition boxes.

[1190] In example embodiments, the method further includes sending the summarized data to one or more data centers.

[1191] In example embodiments, summarizing the data over a given time period to reduce the size of the data includes determining a speed at which data can be sent via a network, wherein the size of the summarized data corresponds to the speed at which data can be sent continuously in real time via the network.

[1192] In example embodiments, the method further includes continuously sending the summarized data to an external device via the network.

[1193] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs and self-organizing at least one of: (i) a storage operation of the data; (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the storage operation includes storing the data in a local database, summarizing the data over a given time period to reduce a size of the data, committing the summarized data to a local ledger, identifying one or more other accessible signal acquisition instruments on an accessible network, and synchronizing the summarized data at the local ledger with at least one of the other accessible signal acquisition instruments. In example embodiments, the method further includes receiving a remote stream of sensor data from one or more other accessible signal acquisition instruments via a network.

[1194] In example embodiments, the method further includes sending an advertisement message to a potential client indicating availability of at least one of the locally stored data, the summarized data, and the remote stream of sensor data.

[1195] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs;

[1196] sampling data received from the sensor inputs, self-organizing at least one of: (i) a storage operation of the data (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the storage operation includes storing the data in a local database, and summarizing the data over a given time period to reduce a size of the data, identifying one or more other accessible signal acquisition instruments on an accessible network, nominating at least one of the one or more other accessible signal acquisition instruments as a logical communication hub, and providing the logical communication hub with a list of available data and their associated sources.

[1197] In example embodiments, the list of available data and their associated sources is provided to the logical communication hub utilizing a hybrid peer-to-peer communications protocol.

[1198] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the storage operation includes storing

the data in a local database, summarizing the data over a given time period to reduce a size of the data, storing the data in a local database, and automatically organizing at least one parameter of the database utilizing machine learning, wherein the organizing is based at least in part on receiving information regarding at least one of an accuracy of classification and an accuracy of prediction of an external machine learning system that uses data from the database.

[1199] In aspects, the collection operation of sensors that provide the plurality of sensor inputs can include receiving instructions directing a mobile data collector unit (e.g., data collector 12008) to operate sensors at a target (e.g., 12002), wherein at least one of the plurality of sensors is arranged in the mobile data collector unit. A communication can be transmitted to one or more other mobile data collector units (12008) regarding the instructions. The swarm 12006 or portion thereof can self-organize a distribution of the mobile data collector unit and the one or more other mobile data collector units (e.g., data collectors 12008) at the target 12002.

[1200] In aspects, self-organizing the distribution of the mobile data collector units at the target 12002 comprises utilizing a machine learning algorithm to determine a respective target location for each of the mobile data collector units. The machine learning algorithm can utilize one or more of a plurality of features to determine the respective target locations. Examples of the features can include: battery life of the mobile data collector units (data collectors 12008), a type of the target 12002 being sensed, a type of signal being sensed, a size of the target 12002, a number of mobile data collector units (data collectors 12008) needed to cover the target 12002, a number of data points needed for the target 12002, a success in prior accomplishment of signal capture, information received from a headquarters or other components from which the instructions are received, and historical information regarding the sensors operated at the target 12002.

[1201] In implementations, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location can include proposing a target location for the mobile data collector unit(s), transmitting the target location to at least one other mobile data collector units, receiving confirmation that there is no contention for the target location, directing one of the mobile data collector units to the target location, and collecting sensor data at the target location from the directed mobile data collector unit.

[1202] Self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location can also include, in certain example embodiments, proposing a target location for the mobile data collector unit, transmitting the target location to at least one of the one or more other mobile data collector units, receiving a proposal for a new target location, directing the mobile data collector unit to the new target location, and collecting sensor data at the new target location from the mobile data collector unit.

[1203] In additional or alternative aspects, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location can comprise proposing a target location for the mobile data collector unit, determining that at least one of the one or more other mobile data collector units is at or moving to the target location, determining a new target location based on the at least one of the one or more other mobile data collector units being at or moving to the target location, directing the mobile data collector unit to

the new target location, and collecting sensor data at the new target location from the mobile data collector unit.

[1204] Self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location can further comprise determining a type of the sensors to operate at the target 12002, receiving confirmation that there is no contention for the type of sensors, directing the mobile data collector unit to operate the type of sensors at the target 12002, and collecting sensor data from the type of sensors at the target 12002 from the mobile data collector unit.

[1205] In aspects, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location can include determining a type of the sensors to operate at the target, transmitting the type of the sensors to at least one of the one or more other mobile data collector units, receiving a proposal for a new type of the sensors, directing the mobile data collector unit to operate the new type of sensors at the target, and collecting sensor data from the new type of sensors at the target from the mobile data collector unit.

[1206] Self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location can include determining a type of the sensors to operate at the target, determining that at least one of the one or more other mobile data collector units is operating or can operate the type of the sensors at the target, determining a new type of the sensors based on the at least one of the one or more other mobile data collector units operating or being capable of operating the type of the sensors at the target, directing the mobile data collector unit to operate the new type of sensors at the target, and collecting sensor data from the new type of sensors at the target from the mobile data collector unit.

[1207] Self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location, in some implementations, can comprise utilizing a swarm optimization algorithm to allocate areas of sensor responsibility amongst the mobile data collector unit and the one or more other mobile data collector units. Examples of the swarm optimization algorithm include but are not limited to Genetic Algorithms (GA), Ant Colony Optimization (ACO), Particle Swarm Optimization (PSO), Differential Evolution (DE), Artificial Bee Colony (ABC), Glowworm Swarm Optimization (GSO), and Cuckoo Search Algorithm (CSA), Genetic Programming (GP), Evolution Strategy (ES), Evolutionary Programming (EP), Firefly Algorithm (FA), Bat Algorithm (BA) and Grey Wolf Optimizer (GWO), or combinations thereof.

[1208] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1209] The present disclosure describes a system for data collection in an industrial environment having automated self-organization, the system according to one disclosed non-limiting

embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the industrial environment and for generating data associated with the plurality of sensor inputs, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1210] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs;

[1211] sampling data received from the sensor inputs and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the collection operation of sensors that provide the plurality of sensor inputs includes receiving instructions directing a mobile data collector unit to operate sensors at a target, wherein at least one of the plurality of sensors is arranged in the mobile data collector unit, transmitting a communication to one or more other mobile data collector units regarding the instructions, and self-organizing a distribution of the mobile data collector unit and the one or more other mobile data collector units at the target.

[1212] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target includes utilizing a machine learning algorithm to determine a respective target location for each of the mobile data collector units.

[1213] In example embodiments, the machine learning algorithm utilizes one or more of a plurality of features to determine the respective target locations, the plurality of features including: battery life of the mobile data collector units, a type of the target being sensed, a type of signal being sensed, a size of the target, a number of mobile data collector units needed to cover the target, a number of data points needed for the target, a success in prior accomplishment of signal capture, information received from a headquarters from which the instructions are received, and historical information regarding the sensors operated at the target.

[1214] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes proposing a target location for the mobile data collector unit, transmitting the target location to at least one of the one or more other mobile data collector units, receiving confirmation that there is no contention for the target location, directing the mobile data collector unit to the target location, and collecting sensor data at the target location from the mobile data collector unit.

[1215] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes proposing a target location for the mobile data collector unit, transmitting the target location to at least one of the one or more other mobile data collector units, receiving a proposal for a new target location, directing the mobile data collector unit to the new target location and collecting sensor data at the new target location from the mobile data collector unit.

[1216] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes proposing a target location for the mobile data collector unit, determining that at least one of the one or more other mobile data collector units is at or moving to the target location, determining a new target location based on the at least one of the one or more other mobile data collector units being at or moving to the target location, directing the mobile data collector unit to the new target location and collecting sensor data at the new target location from the mobile data collector unit.

[1217] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes determining a type of the sensors to operate at the target, receiving confirmation that there is no contention for the type of sensors, directing the mobile data collector unit to operate the type of sensors at the target, and

[1218] collecting sensor data from the type of sensors at the target from the mobile data collector unit.

[1219] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the collection operation of sensors that provide the plurality of sensor inputs includes receiving instructions directing a mobile data collector unit to operate sensors at a target, wherein at least one of the plurality of sensors is arranged in the mobile data collector unit, transmitting a communication to one or more other mobile data collector units regarding the instructions, self-organizing a distribution of the mobile data collector unit and the one or more other mobile data collector units at the target, wherein self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes determining a type of the sensors to operate at the target, transmitting the type of the sensors to at least one of the one or more other mobile data collector units, receiving a proposal for a new type of the sensors, directing the mobile data collector unit to operate the new type of sensors at the target and collecting sensor data from the new type of sensors at the target from the mobile data collector unit.

[1220] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes determining a type of the sensors to operate at the target, determining that at least one of the one or more other mobile data collector units is operating or can operate the type of the sensors at the target, determining a new type of the sensors based on the at least one of the one or more other mobile data collector units operating or being capable of operating the type of the sensors at the target, directing the mobile data collector unit to operate the new type of sensors at the target, and collecting sensor data from the new type of sensors at the target from the mobile data collector unit.

[1221] In example embodiments, self-organizing the distribution of the mobile data collector unit and the one or more other mobile data collector units at the target location includes utilizing a swarm optimization algorithm to allocate areas of sensor responsibility amongst the mobile data collector unit and the one or more other mobile data collector units.

[1222] In example embodiments, the swarm optimization algorithm is one or more types of Genetic Algorithms (GA), Ant Colony Optimization (ACO), Particle Swarm Optimization (PSO), Differential Evolution (DE), Artificial Bee Colony (ABC), Glowworm Swarm Optimization (GSO), and Cuckoo Search Algorithm (CSA), Genetic Programming (GP), Evolution Strategy (ES), Evolutionary Programming (EP), Firefly Algorithm (FA), Bat Algorithm (BA) and Grey Wolf Optimizer (GWO).

[1223] In aspects, the selection operation can comprise receiving a signal relating to at least one condition of the industrial environment 12000 and, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling. The at least one condition of the industrial environment can be a signal-to-noise ratio of the sampled data. The selection operation can include identifying a target signal to be sensed. Additionally, the selection operation further can include identifying one or more non-target signals in a same frequency band as the target signal to be sensed and, based on the identified one or more non-target signals, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1224] The selection operation can comprise identifying other data collectors sensing in a same signal band as the target signal to be sensed, and, based on the identified other data collectors, changing at least one of the sensor inputs analyzed and a frequency of the sampling. In implementations, the selection operation can further comprise identifying a level of activity of a target associated with the target signal to be sensed and, based on the identified level of activity, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1225] The selection operation can further comprise receiving data indicative of environmental conditions near a target associated with the target signal, comparing the received environmental conditions of the target with past environmental conditions near the target or another target similar to the target, and, based on the comparison, changing at least one of the sensor inputs analyzed and a frequency of the sampling. At least a portion of the received sampling data can be transmitted to another data collector according to a predetermined hierarchy of data collection.

[1226] The selection operation further comprises, in some aspects, receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to a quality or sufficiency of the transmitted data, analyzing the received feedback, and, based on the analysis of the received feedback, changing at least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1227] Additionally, or alternatively, the selection operation can comprise receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a

predetermined hierarchy of data collection, receiving feedback via a network connection relating to one or more yield metrics of the transmitted data, analyzing the received feedback, and, based on the analysis of the received feedback, changing at least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1228] In implementations, the selection operation can include receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to power utilization, analyzing the received feedback, and based on the analysis of the received feedback, changing at least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1229] The selection operation can also or alternatively comprise receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to a quality or sufficiency of the transmitted data, analyzing the received feedback, and, based on the analysis of the received feedback, executing a dimensionality reduction algorithm on the sensed data. The dimensionality reduction algorithm can be one or more of a Decision Tree, Random Forest, Principal Component Analysis, Factor Analysis, Linear Discriminant Analysis, Identification based on correlation matrix, Missing Values Ratio, Low Variance Filter, Random Projections, Nonnegative Matrix Factorization, Stacked Auto-encoders, Chi-square or Information Gain, Multidimensional Scaling, Correspondence Analysis, Factor Analysis, Clustering, and Bayesian Models. The dimensionality reduction algorithm can be performed at a data collector 12008, a swarm 12006 of data collectors 12008, a network 12010, a computing system 12012, a data pool 12014, or combination thereof. In aspects, executing the dimensionality reduction algorithm can comprise sending the sensed data to a remote computing device.

[1230] In aspects, a system for self-organizing collection and storage of data collection in a power generation environment can include a data collector for handling a plurality of sensor inputs from various sensors. Such sensors can be a component of the data collector, external to the data collector (e.g., external sensors or components of different data collector(s)), or a combination thereof. The plurality of sensor inputs can be configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system. Examples of such target systems include but are not limited to a fuel handling system, a power source, a turbine, a generator, a gear system, an electrical transmission system, a transformer, a fuel cell, and an energy storage device/system. The system can also include a self-organizing system that can be configured for self-organizing at least one of: (i) a storage operation of the data; (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor input, as is described herein.

[1231] In aspects, the system can include a swarm 12006 of mobile data collectors (e.g., data collectors 12008). Further, in additional or alternative aspects, the self-organizing system can

generate, iterate, optimize, etc. a storage specification for organizing storage of the data. The storage specification, e.g., can specify which data will be stored for local storage in the power generation environment, and which data will be output for streaming via a network connection (e.g., network 12010) from the power generation environment. Other data collection, generation, and/or storage operations can be performed or enabled by the system, as is described herein.

[1232] In a non-limiting example, the system can include a plurality of sensors configured to sense various parameters in the environment of a turbine as a target system. Vibration sensors, temperature sensors, acoustic sensors, strain gauges, and accelerometers, and the like may be utilized by the system to generate data regarding the operation of the turbine. As mentioned herein, any and all of the storage operation, the data collection operation, and the selection operation of the plurality of sensor inputs may be adapted, optimized, learned, or otherwise self-organized by the system.

[1233] In aspects, a system for self-organizing collection and storage of data collection in energy source extraction environment can include a data collector for handling a plurality of sensor inputs from various sensors. Examples of such energy source extraction environments include a coal mining environment, a metal mining environment, a mineral mining environment, and an oil drilling environment, although other extraction environments are contemplated by the present disclosure. The sensors utilized can be a component of the data collector, external to the data collector (e.g., external sensors or components of different data collector(s)), or a combination thereof. The plurality of sensor inputs can be configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system. Examples of such target systems include but are not limited to a hauling system, a lifting system, a drilling system, a mining system, a digging system, a boring system, a material handling system, a conveyor system, a pipeline system, a wastewater treatment system, and a fluid pumping system.

[1234] The system can also include a self-organizing system that can be configured for self-organizing at least one of: (i) a storage operation of the data; (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor input, as is described herein. In aspects, the system can include a swarm 12006 of mobile data collectors (e.g., data collectors 12008). Further, in additional or alternative aspects, the self-organizing system can generate, iterate, optimize, etc. a storage specification for organizing storage of the data. The storage specification, e.g., can specify which data will be stored for local storage in the power generation environment, and which data will be output for streaming via a network connection (e.g., network 12010) from the power generation environment. Other data collection, generation, and/or storage operations can be performed or enabled by the system, as is described herein.

[1235] In a non-limiting example, the system can include a plurality of sensors configured to sense various parameters in the environment of a fluid pumping system as a target system. Vibration sensors, flow sensors, pressure sensors, temperature sensors, acoustic sensors, and the like may be utilized by the system to generate data regarding the operation of the fluid pumping system. As mentioned herein, any and all of the storage operation, the data collection operation, and the

selection operation of the plurality of sensor inputs may be adapted, optimized, learned, or otherwise self-organized by the system.

[1236] In implementations, a system for self-organizing collection and storage of data collection in a manufacturing environment can include a data collector for handling a plurality of sensor inputs from various sensors. Such sensors can be a component of the data collector, external to the data collector (e.g., external sensors or components of different data collector(s)), or a combination thereof. The plurality of sensor inputs can be configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system. Examples of such target systems include but are not limited to a power system, a conveyor system, a generator, an assembly line system, a wafer handling system, a chemical vapor deposition system, an etching system, a printing system, a robotic handling system, a component assembly system, an inspection system, a robotic assembly system, and a semi-conductor production system. The system can also include a self-organizing system that can be configured for self-organizing at least one of: (i) a storage operation of the data; (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor input, as is described herein.

[1237] In aspects, the system can include a swarm 12006 of mobile data collectors (e.g., data collectors 12008). Further, in additional or alternative aspects, the self-organizing system can generate, iterate, optimize, etc. a storage specification for organizing storage of the data. The storage specification, e.g., can specify which data will be stored for local storage in the power generation environment, and which data will be output for streaming via a network connection (e.g., network 12010) from the power generation environment. Other data collection, generation, and/or storage operations can be performed or enabled by the system, as is described herein.

[1238] In a non-limiting example, the system can include a plurality of sensors configured to sense various parameters in the environment of a wafer handling system as a target system. Vibration sensors, fluid flow sensors, pressure sensors, gas sensors, temperature sensors, and the like may be utilized by the system to generate data regarding the operation of the wafer handling system. As mentioned herein, any and all of the storage operation, the data collection operation, and the selection operation of the plurality of sensor inputs may be adapted, optimized, learned, or otherwise self-organized by the system.

[1239] Also disclosed are embodiments of an additional or alternative system for self-organizing collection and storage of data collection in refining environment. Such system(s) can include a data collector for handling a plurality of sensor inputs from various sensors. Examples of such refining environments include a chemical refining environment, a pharmaceutical refining environment, a biological refining environment, and a hydrocarbon refining environment, although other refining environments are contemplated by the present disclosure. The sensors utilized can be a component of the data collector, external to the data collector (e.g., external sensors or components of different data collector(s)), or a combination thereof. The plurality of sensor inputs can be configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system. Examples of such target systems include but are not limited to a power system, a pumping system, a mixing system, a reaction system, a distillation system, a fluid handling system, a heating

system, a cooling system, an evaporation system, a catalytic system, a moving system, and a container system.

[1240] The system can also include a self-organizing system that can be configured for self-organizing at least one of: (i) a storage operation of the data; (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor input, as is described herein. In aspects, the system can include a swarm 12006 of mobile data collectors (e.g., data collectors 12008). Further, in additional or alternative aspects, the self-organizing system can generate, iterate, optimize, etc. a storage specification for organizing storage of the data. The storage specification, e.g., can specify which data will be stored for local storage in the power generation environment, and which data will be output for streaming via a network connection (e.g., network 12010) from the power generation environment. Other data collection, generation, and/or storage operations can be performed or enabled by the system, as is described herein.

[1241] In a non-limiting example, the system can include a plurality of sensors configured to sense various parameters in the refining environment of a heating system as a target system. Temperature sensors, fluid flow sensors, pressure sensors, and the like may be utilized by the system to generate data regarding the operation of the heating system. As mentioned herein, any and all of the storage operation, the data collection operation, and the selection operation of the plurality of sensor inputs may be adapted, optimized, learned, or otherwise self-organized by the system.

[1242] In aspects, a system for self-organizing collection and storage of data collection in a distribution environment can include a data collector for handling a plurality of sensor inputs from various sensors. Such sensors can be a component of the data collector, external to the data collector (e.g., external sensors or components of different data collector(s)), or a combination thereof. The plurality of sensor inputs can be configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system. Examples of such target systems include but are not limited to a power system, a conveyor system, a robotic transport system, a robotic handling system, a packing system, a cold storage system, a hot storage system, a refrigeration system, a vacuum system, a hauling system, a lifting system, an inspection system, and a suspension system. The system can also include a self-organizing system that can be configured for self-organizing at least one of: (i) a storage operation of the data; (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor input, as is described herein.

[1243] In aspects, the system can include a swarm 12006 of mobile data collectors (e.g., data collectors 12008). Further, in additional or alternative aspects, the self-organizing system can generate, iterate, optimize, etc. a storage specification for organizing storage of the data. The storage specification, e.g., can specify which data will be stored for local storage in the power generation environment, and which data will be output for streaming via a network connection (e.g., network 12010) from the power generation environment. Other data collection, generation, and/or storage operations can be performed or enabled by the system, as is described herein.

[1244] In a non-limiting example, the system can include a plurality of sensors configured to sense various parameters in the distribution environment of a refrigeration system as a target system. Power sensors, temperature sensors, vibration sensors, strain gauges, and the like may be utilized by the system to generate data regarding the operation of the turbine. As mentioned herein, any and all of the storage operation, the data collection operation, and the selection operation of the plurality of sensor inputs may be adapted, optimized, learned, or otherwise self-organized by the system.

[1245] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1246] The present disclosure describes a system for data collection in an industrial environment having automated self-organization, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the industrial environment and for generating data associated with the plurality of sensor inputs, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1247] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the selection operation includes

[1248] receiving a signal relating to at least one condition of the industrial environment, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1249] In example embodiments, the at least one condition of the industrial environment is a signal-to-noise ratio of the sampled data.

[1250] In example embodiments, the selection operation includes identifying a target signal to be sensed.

[1251] In example embodiments, the selection operation further includes identifying one or more non-target signals in a same frequency band as the target signal to be sensed, and based on the identified one or more non-target signals, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1252] In example embodiments, the selection operation further includes identifying other data collectors sensing in a same signal band as the target signal to be sensed, and based on the identified other data collectors, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1253] In example embodiments, the selection operation further includes identifying a level of activity of a target associated with the target signal to be sensed, and based on the identified level of activity, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1254] In example embodiments, the selection operation further includes receiving data indicative of environmental conditions near a target associated with the target signal, comparing the received environmental conditions of the target with past environmental conditions near the target or another target similar to the target, and based on the comparison, changing at least one of the sensor inputs analyzed and a frequency of the sampling.

[1255] In example embodiments, the selection operation further includes transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection.

[1256] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the selection operation includes identifying a target signal to be sensed, receiving a signal relating to at least one condition of the industrial environment, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling, receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to a quality or sufficiency of the transmitted data, analyzing the received feedback, and based on the analysis of the received feedback, changing at least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1257] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the selection operation includes identifying a target signal to be sensed, receiving a signal relating to at least one condition of the industrial environment, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling, receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to one or more yield metrics of the transmitted data, analyzing the received feedback, and based on the analysis of the received feedback, changing at

least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1258] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the selection operation includes identifying a target signal to be sensed, receiving a signal relating to at least one condition of the industrial environment, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling, receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback, via a network connection relating to power utilization, analyzing the received feedback, and based on the analysis of the received feedback, changing at least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1259] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the selection operation includes identifying a target signal to be sensed, receiving a signal relating to at least one condition of the industrial environment, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling, receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to a quality or sufficiency of the transmitted data, analyzing the received feedback, and based on the analysis of the received feedback, executing a dimensionality reduction algorithm on the sensed data.

[1260] In example embodiments, the dimensionality reduction algorithm is one or more of a Decision Tree, Random Forest, Principal Component Analysis, Factor Analysis, Linear Discriminant Analysis, Identification based on correlation matrix, Missing Values Ratio, Low Variance Filter, Random Projections, Nonnegative Matrix Factorization, Stacked Auto-encoders, Chi-square or Information Gain, Multidimensional Scaling, Correspondence Analysis, Factor Analysis, Clustering, and Bayesian Models.

[1261] In example embodiments, the dimensionality reduction algorithm is performed at a data collector.

[1262] In example embodiments, executing the dimensionality reduction algorithm includes sending the sensed data to a remote computing device.

[1263] The present disclosure describes a method for data collection in an industrial environment having self-organization functionality, the method according to one disclosed non-limiting embodiment of the present disclosure can include analyzing a plurality of sensor inputs, sampling data received from the sensor inputs, and self-organizing at least one of (i) a storage operation of the data, (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs, wherein the selection operation includes identifying a target signal to be sensed, receiving a signal relating to at least one condition of the industrial environment, based on the signal, changing at least one of the sensor inputs analyzed and a frequency of the sampling, receiving data indicative of environmental conditions near a target associated with the target signal, transmitting at least a portion of the received sampling data to another data collector according to a predetermined hierarchy of data collection, receiving feedback via a network connection relating to at least one of a bandwidth and a quality or of the network connection, analyzing the received feedback, and based on the analysis of the received feedback, changing at least one of the sensor inputs analyzed, the frequency of sampling, the data stored, and the data transmitted.

[1264] The present disclosure describes a system for self-organizing collection and storage of data collection in a power generation environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the power generation environment, wherein the plurality of sensor inputs is configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system selected from a group consisting of a fuel handling system, a power source, a turbine, a generator, a gear system, an electrical transmission system, and a transformer, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1265] In example embodiments, the self-organizing system organizes a swarm of mobile data collectors to collect data from a plurality of target systems.

[1266] In example embodiments, the self-organizing system generates a storage specification for organizing storage of the data, the storage specification specifying data for local storage in the power generation environment and specifying data for streaming via a network connection from the power generation environment.

[1267] The present disclosure describes a system for self-organizing collection and storage of data collection in an energy source extraction environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the energy extraction environment, wherein the plurality of sensor inputs is configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system selected from a group consisting of a hauling system, a lifting system, a drilling system, a mining system, a digging system, a boring system, a material handling system, a conveyor system, a pipeline system, a wastewater treatment system, and a fluid pumping system, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data,

(ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1268] In example embodiments, the self-organizing system organizes a swarm of mobile data collectors to collect data from a plurality of target systems.

[1269] In example embodiments, the self-organizing system generates a storage specification for organizing storage of the data, the storage specification specifying data for local storage in the energy extraction environment and specifying data for streaming via a network connection from the energy extraction environment.

[1270] In example embodiments, the energy source extraction environment is a coal mining environment.

[1271] In example embodiments, the energy source extraction environment is a metal mining environment.

[1272] In example embodiments, the energy source extraction environment is a mineral mining environment.

[1273] In example embodiments, the energy source extraction environment is an oil drilling environment.

[1274] The present disclosure describes a system for self-organizing collection and storage of data collection in a manufacturing environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the power generation environment, wherein the plurality of sensor inputs is configured to sense at least one of an operational mode, a fault mode, and a health status of at least one target system selected from a group consisting of a power system, a conveyor system, a generator, an assembly line system, a wafer handling system, a chemical vapor deposition system, an etching system, a printing system, a robotic handling system, a component assembly system, an inspection system, a robotic assembly system, and a semi-conductor production system, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1275] In example embodiments, the self-organizing system organizes a swarm of mobile data collectors to collect data from a plurality of target systems.

[1276] In example embodiments, the self-organizing system generates a storage specification for organizing the storage of the data, the storage specification specifying data for local storage in the manufacturing environment and specifying data for streaming via a network connection from the manufacturing environment.

[1277] The present disclosure describes a system for self-organizing collection and storage of data collection in a refining environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the power generation environment, wherein the plurality of sensor inputs is configured to sense at least one of an operational mode, a fault mode and a health status of at least one target system selected from a group consisting of a power system, a pumping system, a mixing

system, a reaction system, a distillation system, a fluid handling system, a heating system, a cooling system, an evaporation system, a catalytic system, a moving system, and a container system, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1278] In example embodiments, the self-organizing system organizes a swarm of mobile data collectors to collect data from a plurality of target systems.

[1279] In example embodiments, the self-organizing system generates a storage specification for organizing the storage of the data, the storage specification specifying data for local storage in the refining environment and specifying data for streaming via a network connection from the refining environment.

[1280] In example embodiments, the refining environment is a chemical refining environment.

[1281] In example embodiments, the refining environment is a pharmaceutical refining environment.

[1282] In example embodiments, the refining environment is a biological refining environment.

[1283] In example embodiments, the refining environment is a hydrocarbon refining environment.

[1284] The present disclosure describes a system for self-organizing collection and storage of data collection in a distribution environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include a data collector for handling a plurality of sensor inputs from sensors in the distribution environment, wherein the plurality of sensor inputs is configured to sense at least one of an operational mode, a fault mode and a health status of at least one target system selected from a group consisting of a power system, a conveyor system, a robotic transport system, a robotic handling system, a packing system, a cold storage system, a hot storage system, a refrigeration system, a vacuum system, a hauling system, a lifting system, an inspection system, and a suspension system, and a self-organizing system for self-organizing at least one of (i) a storage operation of the data, (ii) a data collection operation of the sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1285] In example embodiments, the self-organizing system organizes a swarm of mobile data collectors to collect data from a plurality of target systems.

[1286] In example embodiments, the self-organizing system generates a storage specification for organizing the storage of the data, the storage specification specifying data for local storage in the distribution environment and specifying data for streaming via a network connection from the distribution environment.

[1287] Referencing Figure 119, an example system 12200 for self-organized, network-sensitive data collection in an industrial environment is depicted. The system 12200 includes an industrial system 12202 having a number of components 12204, and a number of sensors 12206, wherein each of the sensors 12206 is operatively coupled to at least one of the components 12204. The selection, distribution, type, and communicative setup of sensors depends upon the application of the system 12200 and/or the context.

[1288] In certain example embodiments, sensor data values 12204 are provided to a data collector 12208, which may be in communication with multiple sensors 12206 and/or with a controller 12212. In certain example embodiments, a plant computer 12210 is additionally or alternatively present. In the example system, the controller 12212 is structured to functionally execute operations of the sensor communication circuit 12224, sensor data storage profile circuit 12226, sensor data storage implementation circuit 12228, storage planning circuit 12230, and/or haptic feedback circuit 12232. The controller 12212 is depicted as a separate device for clarity of description. Aspects of the controller 12212 may be present on the sensors 12206, the data controller 12208, the plant computer 12210, and/or on a cloud computing device 12214. In certain embodiments described throughout this disclosure, all aspects of the controller 12212 or other controllers may be present in another device depicted on the system 12200. The plant computer 12210 represents local computing resources, for example processing, memory, and/or network resources, that may be present and/or in communication with the industrial system 12200. In certain example embodiments, the cloud computing device 12214 represents computing resources externally available to the industrial system 12202, for example over a private network, intra-net, through cellular communications, satellite communications, and/or over the internet. In certain example embodiments, the data controller 12208 may be a computing device, a smart sensor, a MUX box, or other data collection device capable to receive data from multiple sensors and to pass-through the data and/or store data for later transmission. An example data controller 12208 has no storage and/or limited storage, and selectively passes sensor data therethrough, with a subset of the sensor data being communicated at a given time due to bandwidth considerations of the data controller 12208, a related network, and/or imposed by environmental constraints. In certain example embodiments, one or more sensors and/or computing devices in the system 12200 are portable devices such as the user associated device 12216 associated with a user 12218, for example a plant operator walking through the industrial system may have a smart phone, which the system 12200 may selectively utilize as a data controller 12208, sensor 12206 – for example to enhance communication throughput, sensor resolution, and/or as a primary method for communicating sensor data values 12244 to the controller 12212. The system 12200 depicts the controller 12212, the sensors 12206, the data controller 12208, the plant computer 12210, and/or the cloud computing device 12214 having a memory storage for storing sensor data thereon, any one or more of which may not have a memory storage for storing sensor data thereon.

[1289] The example system 12200 further includes a mesh network 12220 having a plurality of network nodes depicted thereupon. The mesh network 12220 is depicted in a single location for convenience of illustration, but it will be understood that any network infrastructure that is within the system 12200, and/or within communication with the system 12200, including intermittently, is contemplated within the system network. Additionally, any or all of the cloud server 12214, plant computer 12210, controller 12212, data controller 12208, any network capable sensor 12206, and/or user associated device 12218 may be a part of the network for the system, including a mesh network 12220, during at least certain operating conditions of the system 12200. Additionally, or alternatively, the system 12200 may utilize a hierarchical network, a peer-to-peer network, a peer-

to-peer network with one or more super-nodes, combinations of these, hybrids of these, and/or may include multiple networks within the system 12200 or in communication with the system. It will be appreciated that certain features and operations of the present disclosure are beneficial to only one or more than one of these types of networks, certain features and operations of the present disclosure are beneficial to any type of network, and certain features and operations are particularly beneficial to combinations of these networks, and/or to networks having multiple networking options within the network, where the benefits relate to the utilization of options of any type, or where the benefits relate to one or more options being of a specific network type.

[1290] Referencing Figure 120, an example apparatus 12222 includes the controller 12212 having a sensor communication circuit 12224 that interprets a number of sensor data values 12244 from the number of sensors 12206 and a system collaboration circuit 12228 that communicates at least a portion of the number of sensor data values (e.g., sensor data 12244 to target storage 12252) to a storage target computing device according to a sensor data transmission protocol 12232. The target computing device includes any device in the system having memory that is the target location for the selected sensor data 12252. For example, the cloud server 12214, plant computer 12210, the user associated device 12218, and/or another portion of the controller 12212 that communicates with the sensor 12206 and/or data controller 12208 over the network of the system. The target computing device may be a short-term target (e.g., until a process operation is completed), a medium-term target (e.g., to be held until certain processing operations are completed on the data, and/or until a periodic data migration occurs), and/or a long-term target (e.g., to be held for the course of a data retention policy, and/or until a long-term data migration is planned), and/or the data storage target for an unknown period (e.g., data is passed to a cloud server 12214, whereupon the system 12200, in certain example embodiments, does not maintain control of the data). In certain example embodiments, the target computing device is the next computing device in the system planned to store the data. In certain example embodiments, the target computing device is the next computing device in the system where the data will be moved, where such a move occurs across any aspect of the network of the system 12200.

[1291] The example controller 12212 includes a transmission environment circuit 12226 that determines transmission conditions 12254 corresponding to the communication of the at least a portion of the number of sensor data values 12252 to the storage target computing device. Transmission conditions 12254 include any conditions affecting the transmission of the data. For example, referencing Figure 123, example and non-limiting transmission conditions 12254 are depicted including environmental conditions 12272 (e.g., EM noise, vibration, temperature, the presence and layout of devices or components affecting transmission, such as metal, conductive, or high density) including environmental conditions 12272 that affect communications directly, and environmental conditions 12272 that affect network devices such as routers, servers, transmitters/transceivers, and the like. An example transmission conditions 12254 includes a network performance 12274, such as the specifications of network equipment or nodes, specified limitations of network equipment or nodes (e.g., utilization limits, authorization for usage, available power, etc.), estimated limitations of the network (e.g., based on equipment temperatures,

noise environment, etc.), and/or actual performance of the network (e.g., as observed directly such as by timing messages, sending diagnostic messages, or determining throughput, and/or indirectly by observing parameters such as memory buffers, arriving messages, etc. that tend to provide information about the performance of the network). Another example transmission condition 12254 includes network parameters 12276, such as timing parameters 12278 (e.g., clock speeds, message speeds, synchronous speeds, asynchronous speeds, and the like), protocol selections 12280 (e.g., addressing information, message sizes including administrative support bits within messages, and/or speeds supported by the protocols present or available), file type selections 12282 (e.g., data transfer file types, stored file types, and the network implications such as how much data must be transferred before data is at least partially readable, how to determine data is transferred, likely or supported file sizes, and the like), streaming parameter selections 12284 (e.g., streaming protocols, streaming speeds, priority information of streaming data, available nodes and/or computing devices to manage the streaming data, and the like), and/or compression parameters 12286 (e.g., compression algorithm and type, processing implications at each end of the message, lossy versus lossless compression, how much information must be passed prior to usable data being available, and the like).

[1292] Referencing Figure 124, certain further non-limiting examples of transmission conditions 12254 corresponding to the communication of the sensor data 12252 are depicted. Example and non-limiting transmission conditions 12254 include a mesh network need 12288 (e.g., to rearrange the mesh to balance throughput), a parent node connectivity change 12290 in a hierarchically arranged network (e.g., the parent node has lost connectivity, re-gained connectivity, and/or has changed to a different set of child nodes and/or higher nodes), and/or a network super-node in a hybrid peer-to-peer application-layer network has been replaced 12292. A super-node, as utilized herein, is a node having additional capability from other peer-to-peer nodes. Such additional capability may be by design only – for example a super-node may connect in a different manner and/or to nodes outside of the peer-to-peer node system. In certain example embodiments, the super-node may additionally or alternatively have more processing power, increased network speed or throughput access, and/or more memory (e.g., for buffering, caching, and/or short term storage) to provide more capability to meet the functions of the super-node.

[1293] An example transmission condition 12254 includes a node in a mesh or hierarchical network detected as malicious (e.g., from another supervisory process, heuristically, or as indicated to the system 12200), a peer node has experienced a bandwidth or connectivity change 12296 (e.g., mesh network peer that was forwarding packets has lost connectivity, gained additional bandwidth, had a reduction in available bandwidth, and/or has regained connectivity). An example transmission condition 12254 includes a change in a cost of transmitting information 12298 (e.g., cost has increased or decreased, where cost may be a direct cost parameter such as a data transmission subscription cost, or an abstracted cost parameter reflecting overall system priorities, and/or a current cost of delivering information over a network hop has changed), a change has been made in a hierarchical network arrangement (e.g., network arrangement change 12300) such as to balance bandwidth use in a network tree; and/or a change in a permission scheme 12302 (e.g., a

portion of the network relaying sampling data has had a change in permissions, authorization level, or credentials). Certain further example transmission conditions 12254 include the availability of an additional connection type 12304 (e.g., a higher-bandwidth network connection type has become available, and/or a lower-cost network connection type has become available); a change has been made in a network topology 12306 (e.g., a node has gone offline or online, a mesh change has occurred, and/or a hierarchy change has occurred); and/or a data collection client changed a preference or a requirement 12308 (e.g., a data frequency requirement for at least one of the number of sensor values; a data type requirement for at least one of the number of sensor values; a sensor target for data collection; and/or a data collection client has changed the storage target computing device, which may change the network delivery outcomes and routing).

[1294] The example controller 12212 includes a network management circuit 12230 that updates the sensor data transmission protocol 12232 in response to the transmission conditions 12254. For example, where the transmission conditions 12254 indicate that a current routing, protocol, delivery frequency, delivery rate, and/or any other parameter associated with communicating the sensor data 12252 is no longer cost effective, possible, optimal, and/or where an improvement is available, the network management circuit 12230 updates the sensor data transmission protocol 12232 in response to a lower cost, possible, optimal, and/or improved transmission condition. The example system collaboration circuit 12228 is further responsive to the updated sensor data transmission protocol 12232 – for example, implementing subsequent communications of the sensor data 12252 in compliance with the updated sensor data transmission protocol 12232, providing a communication to the network management circuit 12230 indicating which aspects of the updated sensor data transmission protocol 12232 cannot be or are not being followed, and/or providing an alert (e.g., to an operator, a network node, controller 12212, and/or the network management circuit 12230) indicating that a change is requested, indicating that a change is being implemented, and/or indicating that a requested change cannot be or is not being implemented.

[1295] An example system 12200 includes the transmission conditions 12254 being environmental conditions 12272 relating to sensor communication of the number of sensor data values 12252, where the network management circuit 12230 further analyzes the environmental conditions 12272, and where updating the sensor data transmission protocol 12232 includes modifying the manner in which the number of sensor data values are transmitted from the number of sensors 12206 to the storage target computing device. An example system further includes a data collector 12208 communicatively coupled to at least a portion of the number of sensors 12206 and responsive to the sensor data transmission protocol 12232, where the system collaboration circuit 12228 further receives the number of sensor data values 12244 from the at least a portion of the number of sensors, and where the transmission conditions 12254 correspond to at least one network parameter corresponding to the communication of the number of sensor data values from the at least a portion of the number of sensors. Referencing Figure 125, a number of example sensor data transmission protocol 12232 values are depicted. An example sensor data transmission protocol 12232 value includes a data collection rate 12310 – for example a rate and/or a frequency at which a sensor 12206 transmits, provides, or samples data, and/or at which the data collector 12208

receives, passes along, stores, or otherwise captures sensor data. An example network management circuit 12230 further updates the sensor data transmission protocol 12232 to modify the data collector 12208 to adjust a data collection rate 12310 for at least one of the number of sensors. Another example sensor data transmission protocol 12232 value includes a multiplexing schedule 12312, which includes a data collector 12208 and/or a smart sensor configured to provide multiple sensor data values, such as in an alternating or other scheduled manner, and/or to package multiple sensor values into a single message in a configured manner. An example network management circuit 12230 updates the sensor data transmission protocol 12232 to modify a multiplexing schedule of the data collector 12208 and/or smart sensor. Another example sensor data transmission protocol 12232 value includes an intermediate storage operation 12314, where an intermediate storage is a storage at any location in the system at least one network transmission prior to the target storage computing device. Intermediate storage may be implemented as an on-demand operation, where a request of the data (e.g., from a user, a machine learning operation, or another system component) results in the subsequent transfer from the intermediate storage to the target computing device, and/or the intermediate storage may be implemented to time shift network communications to lower cost and/or lower network utilization times, and/or to manage moment-to-moment traffic on the network. The example network management circuit 12230 updates the sensor data transmission protocol 12232 to command an intermediate storage operation for at least a portion of the number of sensor data values, where the intermediate storage may be on a sensor, data collector, a node in the mesh network, on the controller, on a component, and/or in any other location within the system. An example sensor data transmission protocol 12232 includes a command for further data collection 12316 for at least a portion of the number of sensors – for example because a resolution, rate, and/or frequency of a sensor data provision is not sufficient for some aspect of the system, to provide additional data to a machine learning algorithm, and/or because a prior resource limitation is no longer applicable and further data from one or more sensors is now available. An example sensor data transmission protocol 12232 includes a command to implement a multiplexing schedule 12318 – for example where a data collector 12208 and/or smart sensor is capable to multiplex sensor data but does not do so under all operating conditions, or only does so in response to the multiplexing schedule 12318 of the sensor data transmission protocol 12232.

[1296] An example network management circuit 12230 further updates the sensor data transmission protocol 12232 to adjust a network transmission parameter (e.g., any network parameter 12276) for at least a portion of the number of sensor values. For example, certain network parameters that are not control variables and/or are not currently being controlled are transmission conditions 12254, and certain network parameters are control variables and subject to change in response to the data transmission protocol 12232, and/or the network management circuit 12230 can optionally take control of certain network parameters to make them control variables. An example network management circuit 12230 further updates the sensor data transmission protocol 12232 to change any one or more of: a frequency of data transmitted; a quantity of data transmitted; a destination of data transmitted (including a target or intermediate

destination, and/or a routing); a network protocol used to transmit the data; and/or a network path (e.g., providing a redundant path to transmit the data (e.g., where high noise, high network loss, and/or critical data are involved, the network management circuit 12230 may determine that the system operations are improved with redundant pathing for some of the data)). An example network management circuit 12230 further updates the sensor data transmission protocol 12232, such as to: bond an additional network path to transmit the data (e.g., the network management circuit 12230 may have authority to bring additional network resources online, and/or selectively access additional network resources); re-arrange a hierarchical network to transmit the data (e.g., add or remove a hierarchy layer, change a parent-child relationship, etc., for example, to provide critical data with additional paths, fewer layers, and/or a higher priority path); rebalance a hierarchical network to transmit the data; and/or reconfigure a mesh network to transmit the data. An example network management circuit 12230 further updates the sensor data transmission protocol 12232 to delay a data transmission time, and/or delay the data transmission time to a lower cost transmission time.

[1297] An example network management circuit further updates the sensor data transmission protocol 12232 to reduce the amount of information sent at one time over the network and/or updates the sensor data transmission protocol to adjust a frequency of data sent from a second data collector (e.g., an offset data collector within or not within the direct purview of the network management circuit 12230, but where network resource utilization from the second data collector competes with utilization of the first data collector).

[1298] An example network management circuit 12230 further adjusts an external data access frequency 12234 – for example where the expert system 12242 and/or the machine learning algorithm 12248 access external data 12246 to make continuous improvements to the system (e.g., accessing information outside of the sensor data values 12244, and/or from offset systems or aggregated cloud based data), and/or an external data access timing (12236). The control of external data 12246 access allows for control of network utilization when the system is low on resources, when high fidelity and/or frequency of sensor data values 12244 is prioritized, and/or shifting of resource utilization into lower cost portions of the operating space of the system. In certain example embodiments, the system collaboration circuit 12228 accesses the external data 12246, and is responsive to the adjusted external data access frequency 12234 and/or external data access timing value 12236. An example network management circuit 12230 further adjusts a network utilization value 12238 – for example to keep system utilization operations below a threshold to reserve margin and/or to avoid the need for capital cost upgrades to the system due to capacity limitations. An example network management circuit 12230 adjusts the network utilization value 12238 to utilize bandwidth at a lower cost bandwidth time – for example when competing traffic is lower, when network utilization does not adversely affect other system processes, and/or when power consumption costs are lower.

[1299] An example network management circuit further 12230 enables utilizing a high-speed network, and/or requests a higher cost bandwidth access, for example when system process improvements are sufficient that higher costs are justified, to meet a minimum delivery

requirement for data, and/or to move aging data from the system before it becomes obsolete or must be deleted to make room for subsequent data.

[1300] An example network management circuit 12230 further includes an expert system 12242, where the updating the sensor data transmission protocol 12232 is further in response to operations of the expert system 12242. The self-organized, network-sensitive data collection system may manage or optimize any such parameters or factors noted throughout this disclosure, individually or in combination, using an expert system, which may involve a rule-based optimization, optimization based on a model of performance, and/or optimization using machine learning/artificial intelligence, optionally including deep learning approaches, or a hybrid or combination of the above. Referencing Figure 119, a number of non-limiting examples of expert systems 12242, any one or more of which may be present in embodiments having an expert system 12242. Without limitation to any other aspect of the present disclosure for expert systems, machine learning operations, and/or optimization routines, example expert systems 12242 include a rule-based system 12202 (e.g., seeded by rules based on modeling, expert input, operator experience, or the like); a model-based system 12204 (e.g., modeled responses or relationships in the system informing certain operations of the expert system, and/or working with other operations of the expert system); a neural-net system (e.g., including rules, state machines, decision trees, conditional determinations, and/or any other aspects); a Bayesian-based system 12208 (e.g., statistical modeling, management of probabilistic responses or relationships, and other determinations for managing uncertainty); a fuzzy logic-based system 12210 (e.g., determining fuzzification states for various system parameters, state logic for responses, and de-fuzzification of truth values, and/or other determinations for managing vague states of the system); and/or a machine learning system 12212 (e.g., recursive, iterative, or other long-term optimization or improvement of the expert system, including searching data, resolutions, sampling rates, etc. that are not within the scope of the expert system to determine if improved parameters are available that are not presently utilized), which may be in addition to or an embodiment of the machine learning algorithm 12248. Any aspect of the expert system 12242 may be re-calibrated, deleted, and/or added during operations of the expert system 12242, including in response to updated information learned by the system, provided by a user or operator, provided by the machine learning algorithm 12248, information from external data 12246 and/or from offset systems.

[1301] An example network management circuit 12230 further includes a machine learning algorithm 12248, where updating the sensor data transmission protocol 12232 is further in response to operations of the machine learning algorithm 12248. An example machine learning algorithm 12248 utilizes a machine learning optimization routine, and upon determining that an improved sensor data transmission protocol 12232 is available, the network management circuit 12230 provides the updated sensor data transmission protocol 12232 which is utilized by the system collaboration circuit 12228. In certain example embodiments, the network management circuit 12230 may perform various operations such as supplying a sensor data transmission protocol 12232 which is utilized by the system collaboration circuit 12228 to produce real-world results, applies modeling to the system (either first principles modeling based on system characteristics, a

model utilizing actual operating data for the system, a model utilizing actual operating data for an offset system, and/or combinations of these) to determine what an outcome of a given sensor data transmission protocol 12232 will be or would have been (including, for example, taking extra sensor data beyond what is utilized to support a process operated by the system, and/or utilizing external data 12246 and/or benchmarking data 12240), and/or applying randomized changes to the sensor data transmission protocol 12232 to ensure that an optimization routine does not settle into a local optimum or non-optimal condition.

[1302] An example machine learning algorithm 12248 further utilizes feedback data including the transmission conditions 12254, at least a portion of the number of sensor values 12244; and/or where the feedback data includes benchmarking data 12240. Referencing Figure 126, non-limiting examples of benchmarking data 12240 are depicted. Benchmarking data 12240 may reference, generally, expected data (e.g., according to an expert system 12242, user input, prior experience, and/or modeling outputs), data from an offset system (including as adjusted for differences in the contemplated system 12200), aggregated data for similar systems (e.g., as external data 12246 which may be cloud-based), and the like. Benchmarking data may be relative to the entire system, the network, a node on the network, a data collector, and/or a single sensor or selected group of sensors. Example and non-limiting benchmarking data includes a network efficiency 12320 (e.g., throughput capability, power utilization, quality and/or integrity of communications relative to the infrastructure, load cycle, and/or environmental conditions of the system 12200), a data efficiency 12322 (e.g., a percentage of overall successful data captured relative to a target value, a description of data gaps relative to a target value, and/or may be focused on critical or prioritized data), a comparison with offset data collectors 12324 (e.g., comparing data collectors in the system having a similar environment, data collection responsibility, or other characteristic making the comparison meaningful), a throughput efficiency 12326 (e.g., a utilization of the available throughput, a variability indicator – such as high variability being an indication that a network may be oversized or have further transmission capability, or high variability being an indication that the network is responsive to cost avoidance opportunities – or both depending upon the further context which can be understood looking at other information such as why the utilization differences occur), a data efficacy 12328 (e.g., a determination that captured parameters are result effective, strong control parameters, and/or highly predictive parameters, and that efficacious data is taken at acceptable resolution, sampling rate, and the like), a data quality 12330 (e.g., degradation of the data due to noise, deconvolution errors, multiple calculation operations and rounding, compression, packet losses, etc.), a data precision 12342 (e.g., a determination that sufficiently precise data is taken, preserved during communications, and preserved during storage), a data accuracy 12340 (e.g., a determination that corrupted data, degradation through transmission and/or storage, and/or time lag results in data that is alone inaccurate, or inaccurate as applied in a time sequence or other configuration), a data frequency 12338 (e.g., a determination that data as communicated has sufficient time and/or frequency domain resolution to determine the responses of interest), an environmental response 12336 (e.g., environmental effects on the network are sufficiently managed to maintain other aspects of the data), a signal diversity 12332 (e.g., whether systematic

gaps exist which increase the consequences of degradation – e.g., 1% of the data is missing, but it's s systematically a single critical sensor; do critical sensed parameters have multiple potential sources of information), a critical response (is data sufficient to detect critical responses, such as support for a sensor fusion operation and/or a pattern recognition operation), and/or a mesh networking coherence 12334 (e.g., keeping processors, nodes, and other network aspects together on a single view of applicable memory states).

[1303] Referencing Figure 127, certain further non-limiting examples of benchmarking data 12240 are depicted. Example and non-limiting benchmarking data 12240 includes a data coverage 12346 (e.g., what fraction of the desired data, critical data, etc. was successfully communicated and captured; how is the data distributed throughout the system), a target coverage 12344 (e.g., does a component or process of the system have sufficient time and spatial resolution of sensed values), a motion efficiency 12348 (e.g., reflecting an amount of time, number of steps, or extent of motion required to accomplish a given result, such as where an action is required by a human operator, robotic element, drone, or the like to accomplish an action), a quality of service commitment 12358 (e.g., an agreement, formal or informal commitment, and/or best practice quality of service such as maximum data gaps, minimum up-times, minimum percentages of coverage), a quality of service guarantee 12360 (e.g., a formal agreement to a quality of service with known or modeled consequences that can act in a cost function, etc.), a service level agreement 12362 (e.g., minimum uptimes, data rates, data resolutions, etc., which may be driven by industry practices, regulatory requirements, and/or formal agreements that certain parameters, detection for certain components, or detection for certain processes in the system will meet data delivery requirements in type, resolution, sample rate, etc.), a predetermined quality of service value (e.g., a user-defined value, a policy for the operator of the system, etc.), and/or a network obstruction value 12364. Example and non-limiting network obstruction values 12364 include a network interference value (e.g., environmental noise, traffic on the network, collisions, etc.), a network obstruction value (e.g., a component, operation, and/or object obstructing wireless or wired communication in a region of the network, or over the entire network), and/or an area of impeded network connectivity (e.g., loss of connectivity for any reason, which may be normal at least intermittently during operations, or power loss, movement of objects through the area, movement of a network node through the area (e.g., a smart phone being utilized as a node), etc.). In certain example embodiments, a network obstruction value 12364 may be caused by interference from a component of the system, an interference caused by one or more of the sensors (e.g., due to a fault or failure, or operation outside an expected range), interference caused by a metallic (or other conductive) object, interference caused by a physical obstruction (e.g., a dense object blocking or reducing transparency to wireless transmissions); an attenuated signal caused by a low power condition (e.g., a brown-out, scheduled power reduction, low battery, etc.); and/or an attenuated signal caused by a network traffic demand in a portion of the network (e.g., a node or group of nodes has high traffic demand during operations of the system).

[1304] Yet another example system includes an industrial system including a number of components, and a number of sensors each operatively coupled to at least one of the number of

components; a sensor communication circuit that interprets a number of sensor data values from the number of sensors; a system collaboration circuit that communicates at least a portion of the number of sensor data values over a network having a number of nodes to a storage target computing device according to a sensor data transmission protocol; a transmission environment circuit that determines transmission feedback corresponding to the communication of the at least a portion of the number of sensor data values over the network; and a network management circuit updates the sensor data transmission protocol in response to the transmission feedback. The example system collaboration circuit is further responsive to the updated sensor data transmission protocol.

[1305] Referencing Figure 121, an example apparatus 12256 for self-organized, network-sensitive data collection in an industrial environment for an industrial system having a network with a number of nodes is depicted. In addition to the aspects of apparatus 12222, apparatus 12256 includes the system collaboration circuit 12228 further sending an alert to at least one of the number of nodes (e.g., as a node communication 12258) in response to the updated sensor data transmission protocol 12232. In certain example embodiments, updating the sensor data transmission protocol 12232 includes the network management circuit 12230 including node control instructions, such as providing instructions to rearrange a mesh network including the number of nodes, providing instructions to rearrange a hierarchical data network including the number of nodes, rearranging a peer-to-peer data network including the number of nodes, rearranging a hybrid peer-to-peer data network including the number of nodes. In certain example embodiments, the system collaboration circuit 12228 provides node control instructions as one or more node communications 12258.

[1306] In certain example embodiments, updating the sensor data transmission protocol 12232 includes the network management circuit 12230 providing instructions to reduce a quantity of data sent over the network; providing instructions to adjust a frequency of data capture sent over the network; providing instructions to time-shift delivery of at least a portion of the number of sensor values sent over the network (e.g., utilizing intermediate storage); providing instructions to change a network protocol corresponding to the network; providing instructions to reduce a throughput of at least one device coupled to the network; providing instructions to reduce a bandwidth use of the network; providing instructions to compress data corresponding to at least a portion of the number of sensor values sent over the network; providing instructions to condense data corresponding to at least a portion of the number of sensor values sent over the network (e.g., providing a relevant subset, reduced sample rate data, etc.); providing instructions to summarize data (e.g., providing a statistical description, an aggregated value, etc.) corresponding to at least a portion of the number of sensor values sent over the network; providing instructions to encrypt data corresponding to at least a portion of the number of sensor values sent over the network (e.g., to enable using an alternate, less secure network path, and/or to access another network path requiring encryption); providing instructions to deliver data corresponding to at least a portion of the number of sensor values to a distributed ledger; providing instructions to deliver data corresponding to at least a portion of the number of sensor values to a central server (e.g., the plant computer 12210 and/or cloud server 12214); providing instructions to deliver data corresponding to at least a portion of

the number of sensor values to a super-node; and providing instructions to deliver data corresponding to at least a portion of the number of sensor values redundantly across a number of network connections. In certain example embodiments, updating the sensor data transmission includes providing instructions to deliver data corresponding to at least a portion of the number of sensor values to one of the components (e.g., where one or more components 12204 in the system has memory storage and is communicatively accessible to the sensor 12206, the data collector 12208, and/or the network), and/or where the one of the components is communicatively coupled to the sensor providing the data corresponding to at least a portion of the number of sensor values (e.g., where the data to be stored on the component 12204 is the component the data was measured for, or is in proximity to the sensor 12206 taking the data).

[1307] An example network includes a mesh network where the network management circuit 12230 further updates the sensor data transmission protocol 12232 to provide instructions to eject (e.g., remove from the mesh map, take it out of service, etc.) one of the number of nodes from the mesh network. An example network includes a peer-to-peer network, where the network management circuit 12230 further updates the sensor data transmission protocol 12232 to provide instructions to eject one of the number of nodes from the peer-to-peer network.

[1308] An example network management circuit 12230 further updates the sensor data transmission protocol 12232 to cache (e.g., as a sensor data cache 12260) at least a portion of the number of sensor values 12252. In certain further embodiments, the network management circuit 12230 further updates the sensor data transmission protocol 12232 to communicate the cached sensor values 12260 in response to at least one of: a determination that the cached data is requested (e.g., a user, model, machine learning algorithm, expert system, etc. has requested the data); a determination that the network feedback indicates communication of the cached data is available (e.g., a prior limitation on the network leading the network management circuit 12230 to direct caching is now lifted or improved); and/or a determination that higher priority data is present that requires utilization of cache resources holding the cached data 12260.

[1309] An example system 12200 for self-organized, network-sensitive data collection in an industrial environment includes an industrial system 12202 having a number of components 12204 and a number of sensors 12206 each operatively coupled to at least one of the number of components 12204. A sensor communication circuit 12224 interprets the number of sensor data values 12244 from the number of sensors at a predetermined frequency. The system collaboration circuit 12228 that communicates at least a portion of the number of sensor data values 12252 over a network having a number of nodes to a storage target computing device according to the sensor data transmission protocol 12232, where the sensor data transmission protocol 12232 includes a predetermined hierarchy of data collection and the predetermined frequency. An example data management circuit 12230 adjusts the predetermined frequency in response to transmission conditions 12254, and/or in response to benchmarking data 12240.

[1310] An example system 12200 for self-organized, network-sensitive data collection in an industrial environment includes an industrial system 12202 having a number of components 12204, and a number of sensors 12206 each operatively coupled to at least one of the number of

components 12204. The sensor communication circuit 12224 interprets a number of sensor data values 12244 from the number of sensors 12206 at a predetermined frequency, and the system collaboration circuit 12228 communicates at least a portion of the number of sensor data values 12252 over a network having a number of nodes to a storage target computing device according to a sensor data transmission protocol. A transmission environment circuit 12226 determines transmission feedback (e.g., transmission conditions 12254) corresponding to the communication of the at least a portion of the number of sensor data values 12252 over the network. A network management circuit 12230 updates the sensor data transmission protocol 12232 in response to the transmission feedback 12254, and a network notification circuit 12268 provides an alert value 12264 in response to the updated sensor data transmission protocol 12232. Example alert values 12264 include a notification to an operator, a notification to a user, a notification to a portable device associated with a user, a notification to a node of the network, a notification to a cloud computing device, a notification to a plant computing device, and/or a provision of the alert as external data to an offset system. Example and non-limiting alert conditions include a component of the system operating in a fault condition, a process of the system operating in a fault condition, a commencement of the utilization of cache storage and/or intermediate storage for sensor values due to a network communication limit, a change in the sensor data transmission protocol (including changes of a selected type), and/or a change in the sensor data transmission protocol that may result in loss of data fidelity or resolution (e.g., compression of data, condensing of data, and/or summarizing data).

[1311] An example transmission feedback includes a feedback value such as: a change in transmission pricing, a change in storage pricing, a loss of connectivity, a reduction of bandwidth, a change in connectivity, a change in network availability, a change in network range, a change in wide area network (WAN) connectivity, and/or a change in wireless local area network (WLAN) connectivity.

[1312] An example system includes an assembly line industrial system having a number of vibrating components, such as motors, conveyors, fans, and/or compressors. The system includes a number of sensors that determine various parameters related to the vibrating components, including determination of diagnostic and/or process related information (proper operation, off-nominal operation, operating speed, imminent servicing or failure, etc.) of one or more of the components. Example sensors, without limitation, include noise, vibration, acceleration, temperature, and/or shaft speed sensors. The sensor information is conveyed to a target storage system, including at least partially through a network communicatively coupled to the assembly line industrial system. The example system includes a network management circuit that determines a sensor data transmission protocol to control flow of data from the sensors to the target storage system. The network management circuit, a related expert system, and/or a related machine learning algorithm, updates the sensor data transmission protocol to ensure efficient network utilization, sufficient delivery of data to support system control, diagnostics, and/or other determinations planned for the data outside of the system, to reduce resource utilization of data transmission, and/or to respond to system noise factors, variability, and/or changes in the system

or related aspects such as cost or environment parameters. The example system includes improvement of system operations to ensure that diagnostics, controls, or other data dependent operations can be completed, to reduce costs while maintaining performance, and/or to increase system capability over time or process cycles.

[1313] An example system includes an automated robotic handling system, including a number of components such as actuators, gear boxes, and/or rail guides. The system includes a number of sensors that determine various parameters related to the components, including without limitation actuator position and/or feedback sensors, vibration, acceleration, temperature, imaging sensors, and/or spatial position sensors (e.g., within the handling system, a related plant, and/or GPS-type positioning). The sensor information is conveyed to a target storage system, including at least partially through a network communicatively coupled to the automated robotic handling system. The example system includes a network management circuit that determines a sensor data transmission protocol to control flow of data from the sensors to the target storage system. The network management circuit, a related expert system, and/or a related machine learning algorithm, updates the sensor data transmission protocol to ensure efficient network utilization, sufficient delivery of data to support system control, diagnostics, improvement and/or efficiency updates to handling efficiency, and/or other determinations planned for the data outside of the system, to reduce resource utilization of data transmission, and/or to respond to system noise factors, variability, and/or changes in the system or related aspects such as cost or environment parameters. The example system includes improvement of system operations to ensure that diagnostics, controls, or other data dependent operations can be completed, to reduce costs while maintaining performance, and/or to increase system capability over time or process cycles.

[1314] An example system includes a mining operation, including a surface and/or underground mining operation. The example mining operation includes components such as an underground inspection system, pumps, ventilation, generators and/or power generation, gas composition or quality systems, and/or process stream composition systems (e.g., including determination of desired material compositions, and/or composition of effluent streams for pollution and/or regulatory control). Various sensors are present in an example system to support control of the operation, determine status of the components, support safe operation, and/or to support regulatory compliance. The sensor information is conveyed to a target storage system, including at least partially through a network communicatively coupled to the mining operation. In certain example embodiments, the network infrastructure of the mining operation exhibits high variability, due to, without limitation, significant environmental variability (e.g., pit or shaft condition variability) and/or intermittent availability – e.g., shutting off electronics during certain mining operations, difficulty in providing network access to portions of the mining operation, and/or the desirability to include mobile or intermittently available devices within the network infrastructure. The example system includes a network management circuit that determines a sensor data transmission protocol to control flow of data from the sensors to the target storage system. The network management circuit, a related expert system, and/or a related machine learning algorithm, updates the sensor data transmission protocol to ensure efficient network utilization, sufficient delivery of

data to support system control, diagnostics, improvement and/or efficiency updates to handling efficiency, support for financial and/or regulatory compliance, and/or other determinations planned for the data outside of the system, to reduce resource utilization of data transmission, and/or to respond to system noise factors, variability, network infrastructure challenges, and/or changes in the system or related aspects such as cost or environment parameters.

[1315] An example system includes an aerospace system, such as a plane, helicopter, satellite, space vehicle or launcher, orbital platform, and/or missile. Aerospace systems have numerous systems supported by sensors, such as engine operations, control surface status and vibrations, environmental status (internal and external), and telemetry support. Additionally, aerospace systems have high variability in both the number of sensors of varying types (e.g., a small number of fuel pressure sensors, but a large number of control surface sensors) as well as the sampling rates for relevant determinations of sensors of varying types (e.g., 1-second data may be sufficient for internal cabin pressure, but weather radar or engine speed sensors may require much higher time resolution). Computing power on an aerospace application is at a premium due to power consumption and weight considerations, and accordingly iterative, recursive, deep learning, expert system, and/or machine learning operations to improve any systems on the aerospace system, including sensor data taking and transmission of sensor information, are driven in many embodiments to computing devices outside of the aerospace vehicle of the system (e.g., through offline learning, post-processing, or the like). Storage capacity on an aerospace application is similarly at a premium, such that long-term storage of sensor data on the aerospace vehicle is not a cost-effective solution for many embodiments. Additionally, network communication from an aerospace vehicle may be subject to high variability and/or bandwidth limitations as the vehicle moves rapidly through the environment and/or into areas where direct communication with ground-based resources is not practical. Further, certain aerospace applications have significant competition for available network resources – for example in environments with a large number of passengers where passenger utilization of a network infrastructure consumes significant bandwidth. Accordingly, it can be seen that operations of a network management circuit, a related expert system, and/or a related machine learning algorithm, to update the sensor data transmission protocol can significantly enhance sensing operations in various aerospace systems. Additionally, certain aerospace applications have a high number of offset systems, enhancing the ability of an expert system or machine learning algorithm to improve sensor data capture and transmission operations, and/or to manage the high variability in sensed parameters (frequency, data rate, and/or data resolution) for the system across operating conditions.

[1316] An example system includes an oil or gas production system, such as a production platform (onshore or offshore), pumps, rigs, drilling equipment, blenders, and the like. Oil and gas production systems exhibit high variability in sensed variable types and sensing parameters, such as vibration (e.g., pumps, rotating shafts, fluid flow through pipes, etc. – which may be high frequency or low frequency), gas composition (e.g., of a wellhead area, personnel zone, near storage tanks, etc. – where low frequency may typically be acceptable, and/or it may be acceptable that no data is taken during certain times such as when personnel are not present), and/or pressure

values (which may vary significantly both in required resolution and frequency or sampling rate depending upon operations currently occurring in the system). Additionally, oil and gas production systems have high variability in network infrastructure, both according to the system (e.g., an offshore platform versus a long-term ground-based production facility) and according to the operations being performed by the system (e.g., a wellhead in production may have limited network access, while a drilling or fracturing operation may have significant network infrastructure at a site during operations). Accordingly, it can be seen that operations of a network management circuit, a related expert system, and/or a related machine learning algorithm, to update the sensor data transmission protocol can significantly enhance sensing operations in various oil or gas production systems.

[1317] The present disclosure describes system for self-organized, network-sensitive data collection in an industrial environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include an industrial system including a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components, a sensor communication circuit structured to interpret a plurality of sensor data values from the plurality of sensors, a system collaboration circuit structured to communicate at least a portion of the plurality of sensor data values to a storage target computing device according to a sensor data transmission protocol, a transmission environment circuit structured to determine transmission conditions corresponding to the communication of the at least a portion of the plurality of sensor data values to the storage target computing device, a network management circuit structured to update the sensor data transmission protocol in response to the transmission conditions, and wherein the system collaboration circuit is further responsive to the updated sensor data transmission protocol.

[1318] In example embodiments, the transmission conditions include environmental conditions relating to sensor communication of the plurality of sensor data values, and wherein the network management circuit is further structured to analyze the environmental conditions, and wherein updating the sensor data transmission protocol includes modifying the manner in which the plurality of sensor data values is transmitted from the plurality of sensors to the storage target computing device.

[1319] In example embodiments, a data collector communicatively coupled to at least a portion of the plurality of sensors and responsive to the sensor data transmission protocol, wherein the system collaboration circuit is structured to receive the plurality of sensor data values from the at least a portion of the plurality of sensors, and wherein the transmission conditions correspond to at least one network parameter corresponding to the communication of the plurality of sensor data values from the at least a portion of the plurality of sensors.

[1320] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to modify the data collector to adjust a data collection rate for at least one of the plurality of sensors.

[1321] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to modify a multiplexing schedule of the data collector.

[1322] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to command an intermediate storage operation for at least a portion of the plurality of sensor data values.

[1323] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to command further data collection for at least a portion of the plurality of sensors.

[1324] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to modify the data collector to implement a multiplexing schedule.

[1325] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to adjust a network transmission parameter for at least a portion of the plurality of sensor values.

[1326] In example embodiments, the adjusted network transmission parameter includes at least one parameter selected from the parameters consisting of a timing parameter, a protocol selection, a file type selection, a streaming parameter selection, and a compression parameter.

[1327] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to change a frequency of data transmitted.

[1328] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to change a quantity of data transmitted.

[1329] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to change a destination of data transmitted.

[1330] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to change a network protocol used to transmit the data.

[1331] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to add a redundant network path to transmit the data.

[1332] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to bond an additional network path to transmit the data.

[1333] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to re-arrange a hierarchical network to transmit the data.

[1334] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to rebalance a hierarchical network to transmit the data.

[1335] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to reconfigure a mesh network to transmit the data.

[1336] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to delay a data transmission time.

[1337] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to delay the data transmission time to a lower cost transmission time.

[1338] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to reduce the amount of information sent at one time over the network.

[1339] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to adjust a frequency of data sent from a second data collector.

[1340] In example embodiments, the network management circuit is further structured to adjust an external data access frequency, and wherein the system collaboration circuit is responsive to the adjusted external data access frequency.

[1341] In example embodiments, the network management circuit is further structured to adjust an external data access timing value, and wherein the system collaboration circuit is responsive to the adjusted external data access timing value.

[1342] In example embodiments, the network management circuit is further structured to adjust a network utilization value.

[1343] In example embodiments, the network management circuit is further structured to adjust the network utilization value to utilize bandwidth at a lower cost bandwidth time.

[1344] In example embodiments, the network management circuit is further structured to enable utilizing a high-speed network.

[1345] In example embodiments, the network management circuit is further structured to request a higher cost bandwidth access, and to update the sensor transmission protocol in response to the higher cost bandwidth access.

[1346] In example embodiments, the network management circuit further includes an expert system, and wherein the updating the sensor data transmission protocol is further in response to operations of the expert system.

[1347] In example embodiments, the network management circuit further includes a machine learning algorithm, and wherein the updating the sensor data transmission protocol is further in response to operations of the machine learning algorithm.

[1348] In example embodiments, the machine learning algorithm is further structured to utilize feedback data including the transmission conditions.

[1349] In example embodiments, the feedback data further includes at least a portion of the plurality of sensor values.

[1350] In example embodiments, the feedback data further includes benchmarking data.

[1351] In example embodiments, the benchmarking data further includes data selected from the list consisting of: a network efficiency, a data efficiency, a comparison with offset data collectors, a throughput efficiency, a data efficacy, a data quality, a data precision, a data accuracy, and a data frequency.

[1352] In example embodiments, the benchmarking data further includes data selected from the list consisting of: an environmental response, a mesh networking coherence, a data coverage, a target coverage, a signal diversity, a critical response, and a motion efficiency.

[1353] In example embodiments, the transmission conditions corresponding to the communication comprise at least one condition selected from the conditions consisting of a mesh network needs to rearrange to balance throughput, a parent node in a hierarchically arranged network has had a change in connectivity, a network super-node in a hybrid peer-to-peer application-layer network has been replaced, and a node in a mesh or hierarchical network has been detected as malicious.

[1354] In example embodiments, the transmission conditions corresponding to the communication comprise at least one condition selected from the conditions consisting of a mesh network peer forwarding packets has lost connectivity, a mesh network peer forwarding packets has gained additional bandwidth, a mesh network peer forwarding packets has had a reduction in bandwidth, and a mesh network peer forwarding packets has regained connectivity.

[1355] In example embodiments, the transmission conditions corresponding to the communication comprise at least one condition selected from the conditions consisting of a cost of transmitting information has changed dynamically, a change has been made in a hierarchical network arrangement to balance bandwidth use in a network tree, a portion of the network relaying sampling data has had a change in permissions, authorization level, or credentials, a current cost of delivering information over a network hop has changed, a higher-bandwidth network connection type has become available, a lower-cost network connection type has become available, and a change has been made in a network topology.

[1356] In example embodiments, the transmission conditions corresponding to the communication include at least one condition selected from the conditions consisting of a data collection client has changed a data frequency requirement for at least one of the plurality of sensor values, a data collection client has changed a data type requirement for at least one of the plurality of sensor values, a data collection client has changed a sensor target for data collection, and a data collection client has changed the storage target computing device.

[1357] The present disclosure describes system for self-organized, network-sensitive data collection in an industrial environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include an industrial system including a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components, a sensor communication circuit structured to interpret a plurality of sensor data values from the plurality of sensors, a system collaboration circuit structured to communicate at least a portion of the plurality of sensor data values over a network having a plurality of nodes to a storage target computing device according to a sensor data transmission protocol, a transmission environment circuit structured to determine transmission feedback corresponding to the communication of the at least a portion of the plurality of sensor data values over the network, and a network management circuit structured to update the sensor data transmission protocol in response to the transmission feedback, wherein the system collaboration circuit is further responsive to the updated sensor data transmission protocol.

[1358] In example embodiments, the system collaboration circuit is further structured to send an alert to at least one of the plurality of nodes in response to the updated sensor data transmission protocol.

[1359] In example embodiments, updating the sensor data transmission includes at least one operation selected from the operations consisting of providing instructions to rearrange a mesh network including the plurality of nodes, providing instructions to rearrange a hierarchical data network including the plurality of nodes, rearranging a peer-to-peer data network including the plurality of nodes and rearranging a hybrid peer-to-peer data network including the plurality of nodes.

[1360] In example embodiments, updating the sensor data transmission includes at least one operation selected from the operations consisting of providing instructions to reduce a quantity of data sent over the network, providing instructions to adjust a frequency of data capture sent over the network, providing instructions to time-shift delivery of at least a portion of the plurality of sensor values sent over the network, and providing instructions to change a network protocol corresponding to the network.

[1361] In example embodiments, updating the sensor data transmission includes at least one operation selected from the operations consisting of providing instructions to reduce a throughput of at least one device coupled to the network, providing instructions to reduce a bandwidth use of the network, providing instructions to compress data corresponding to at least a portion of the plurality of sensor values sent over the network, providing instructions to condense data corresponding to at least a portion of the plurality of sensor values sent over the network, providing instructions to summarize data corresponding to at least a portion of the plurality of sensor values sent over the network, and providing instructions to encrypt data corresponding to at least a portion of the plurality of sensor values sent over the network.

[1362] In example embodiments, updating the sensor data transmission includes at least one operation selected from the operations consisting of providing instructions to deliver data corresponding to at least a portion of the plurality of sensor values to a distributed ledger, providing instructions to deliver data corresponding to at least a portion of the plurality of sensor values to a central server, providing instructions to deliver data corresponding to at least a portion of the plurality of sensor values to a super-node and providing instructions to deliver data corresponding to at least a portion of the plurality of sensor values redundantly across a plurality of network connections.

[1363] In example embodiments, updating the sensor data transmission includes providing instructions to deliver data corresponding to at least a portion of the plurality of sensor values to one of the plurality of components.

[1364] In example embodiments, the one of the plurality of components is communicatively coupled to the sensor providing the data corresponding to at least a portion of the plurality of sensor values.

[1365] In example embodiments, the system collaboration circuit is further structured to interpret a quality of service commitment, and wherein the network management circuit is further structured to update the sensor data transmission protocol further in response to the quality of service commitment.

[1366] In example embodiments, the system collaboration circuit is further structured to interpret a service level agreement, and wherein the network management circuit is further structured to update the sensor data transmission protocol further in response to the service level agreement.

[1367] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to provide instructions to increase a quality of service value.

[1368] In example embodiments, the network includes a mesh network, and wherein the network management circuit is further structured to update the sensor data transmission protocol to provide instructions to eject one of the plurality of nodes from the mesh network.

[1369] In example embodiments, the network includes a peer-to-peer network, and wherein the network management circuit is further structured to update the sensor data transmission protocol to provide instructions to eject one of the plurality of nodes from the peer-to-peer network.

[1370] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to cache at least a portion of the plurality of sensor values.

[1371] In example embodiments, the network management circuit is further structured to update the sensor data transmission protocol to communicate the cached at least a portion of the plurality of sensor values in response to at least one of a determination that the cached data is requested, a determination that the network feedback indicates communication of the cached data is available, and a determination that higher priority data is present that requires utilization of cache resources holding the cached data.

[1372] In example embodiments, the system further includes a data collector configured to receive the at least a portion of the plurality of sensor data values, wherein the at least a portion of the plurality of sensor data values includes data provided by a plurality of the sensors, and wherein the transmission feedback includes network performance information corresponding to the data collector.

[1373] In example embodiments, the system further includes a data collector configured to receive the at least a portion of the plurality of sensor data values, wherein the at least a portion of the plurality of sensor data values includes data provided by a plurality of the sensors, a second data collector communicatively coupled to the network, and wherein the transmission feedback includes network performance information corresponding to the second data collector.

[1374] The present disclosure describes system for self-organized, network-sensitive data collection in an industrial environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include an industrial system including a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components, a sensor communication circuit structured to interpret a plurality of sensor data values from the plurality of sensors at a predetermined frequency, a system collaboration circuit structured to communicate at least a portion of the plurality of sensor data values over a network having a plurality of nodes to a storage target computing device according to a sensor data transmission protocol, the sensor data transmission protocol including a predetermined hierarchy of data collection and the predetermined frequency, a transmission environment circuit structured to determine transmission feedback corresponding to the communication of the at least a portion of

the plurality of sensor data values over the network, and a network management circuit structured to update the sensor data transmission protocol in response to the transmission feedback and further in response to benchmarking data, wherein the system collaboration circuit is further responsive to the updated sensor data transmission protocol.

[1375] In example embodiments, updating the sensor data transmission includes at least one operation selected from the operations consisting of providing an instruction to change the sensors of the plurality of sensors, providing an instruction to adjust the predetermined frequency, providing an instruction to adjust a quantity of the plurality of sensor data values that are stored, providing an instruction to adjust a data transmission rate of the communication of the at least a portion of the plurality of sensor data values, providing an instruction to adjust a data transmission time of the communication of the at least a portion of the plurality of sensor data values, and providing an instruction to adjust a networking method of the communication over the network.

[1376] In example embodiments, the benchmarking data further includes data selected from the list consisting of a network efficiency, a data efficiency, a comparison with offset data collectors, a throughput efficiency, a data efficacy, a data quality, a data precision, a data accuracy, and a data frequency.

[1377] In example embodiments, the benchmarking data further includes data selected from the list consisting of an environmental response, a mesh networking coherence, a data coverage, a target coverage, a signal diversity, a critical response, and a motion In example embodiments, the benchmarking data further includes data selected from the list consisting of a quality of service commitment, a quality of service guarantee, a service level agreement, and a predetermined quality of service value.

[1378] In example embodiments, the benchmarking data further includes data selected from the list consisting of a network interference value, a network obstruction value, and an area of impeded network connectivity.

[1379] In example embodiments, the transmission feedback includes a communication interference value selected from the values consisting of an interference caused by a component of the system, an interference caused by one of the sensors, an interference caused by a metallic object, an interference caused by a physical obstruction, an attenuated signal caused by a low power condition, and an attenuated signal caused by a network traffic demand in a portion of the network.

[1380] The present disclosure describes a system for self-organized, network-sensitive data collection in an industrial environment, the system according to one disclosed non-limiting embodiment of the present disclosure can include an industrial system including a plurality of components, and a plurality of sensors each operatively coupled to at least one of the plurality of components, a sensor communication circuit structured to interpret a plurality of sensor data values from the plurality of sensors at a predetermined frequency, a system collaboration circuit structured to communicate at least a portion of the plurality of sensor data values over a network having a plurality of nodes to a storage target computing device according to a sensor data transmission protocol, a transmission environment circuit structured to determine transmission feedback corresponding to the communication of the at least a portion of the plurality of sensor data values

over the network, a network management circuit structured to update the sensor data transmission protocol in response to the transmission feedback and a network notification circuit structured to provide an alert value in response to the updated sensor data transmission protocol, wherein the system collaboration circuit is further responsive to the updated sensor data transmission protocol.

[1381] In example embodiments, the transmission feedback includes at least one feedback value selected from the values consisting of: a change in transmission pricing, a change in storage pricing, a loss of connectivity, a reduction of bandwidth, a change in connectivity, a change in network availability, a change in network range, a change in wide area network (WAN) connectivity, and a change in wireless local area network (WLAN) connectivity.

[1382] In example embodiments, the network management circuit further includes an expert system, and wherein the updating the sensor data transmission protocol is further in response to operations of the expert system.

[1383] In example embodiments, the expert system includes at least one system selected from the systems consisting of: a rule-based system, a model-based system, a neural-net system, a Bayesian-based system, a fuzzy logic-based system, and a machine learning system.

[1384] In example embodiments, the network management circuit further includes a machine learning algorithm, and wherein the updating the sensor data transmission protocol is further in response to operations of the machine learning algorithm.

[1385] In example embodiments, the machine learning algorithm is further structured to utilize feedback data including the transmission conditions.

[1386] In example embodiments, the feedback data further includes at least a portion of the plurality of sensor values.

[1387] In example embodiments, the feedback data further includes benchmarking data.

[1388] In example embodiments, the benchmarking data further includes data selected from the list consisting of: a network efficiency, a data efficiency, a comparison with offset data collectors, a throughput efficiency, a data efficacy, a data quality, a data precision, a data accuracy, and a data frequency.

[1389] In example embodiments, the benchmarking data further includes data selected from the list consisting of: an environmental response, a mesh networking coherence, a data coverage, a target coverage, a signal diversity, a critical response, and a motion efficiency.

[1390] Referencing Figure 128, an example system 12500 for data collection in an industrial environment includes an industrial system 12502 having a number of components 12504, and a number of sensors 12506, wherein each of the sensors 12506 is operatively coupled to at least one of the components 12504. The selection, distribution, type, and communicative setup of sensors depends upon the application of the system 12500 and/or the context.

[1391] The example system 12500 further includes a sensor communication circuit 12522 (reference Figure 129) that interprets a number of sensor data values 12542. An example system includes the sensor data values 12542 being a number of values to support a sensor fusion operation, for example a set of sensors believed to encompass detection of operating conditions of the system that affect a desired output, to control a process or portion of the industrial system

12502, to diagnose or predict an aspect of the industrial system 12502 or a process associated with the industrial system industrial system 12502.

[1392] In certain example embodiments, sensor data values 12542 are provided to a data collector 12508, which may be in communication with multiple sensors 12506 and/or with a controller 12512. In certain example embodiments, a plant computer 12510 is additionally or alternatively present. In the example system, the controller 12512 is structured to functionally execute operations of the sensor communication circuit 12522, sensor data storage profile circuit 12524, sensor data storage implementation circuit 12526, storage planning circuit 12528, and/or haptic feedback circuit 12530. The controller 12512 is depicted as a separate device for clarity of description. Aspects of the controller 12512 may be present on the sensors 12506, the data controller 12508, the plant computer 12510, and/or on a cloud computing device 12514. In certain embodiments described throughout this disclosure, all aspects of the controller 12512 or other controllers may be present in another device depicted on the system 12500. The plant computer 12510 represents local computing resources, for example processing, memory, and/or network resources, that may be present and/or in communication with the industrial system 12500. In certain example embodiments, the cloud computing device 12514 represents computing resources externally available to the industrial system 12502, for example over a private network, intra-net, through cellular communications, satellite communications, and/or over the internet. In certain example embodiments, the data controller 12508 may be a computing device, a smart sensor, a MUX box, or other data collection device capable to receive data from multiple sensors and to pass-through the data and/or store data for later transmission. An example data controller 12508 has no storage and/or limited storage, and selectively passes sensor data therethrough, with a subset of the sensor data being communicated at a given time due to bandwidth considerations of the data controller 12508, a related network, and/or imposed by environmental constraints. In certain example embodiments, one or more sensors and/or computing devices in the system 12500 are portable devices – for example a plant operator walking through the industrial system may have a smart phone, which the system 12500 may selectively utilize as a data controller 12508, sensor 12506 – for example to enhance communication throughput, sensor resolution, and/or as a primary method for communicating sensor data values 12542 to the controller 12512. The system 12500 depicts the controller 12512, the sensors 12506, the data controller 12508, the plant computer 12510, and/or the cloud computing device 12514 having a memory storage for storing sensor data thereon, any one or more of which may not have a memory storage for storing sensor data thereon. In certain example embodiments, the sensor data storage profile circuit 12524 prepares a data storage profile 12532 that directs sensor data to memory storage, including moving sensor data in a controlled manner from one memory storage to another. Sensor data stored on various devices consumes memory on the device, transferring the stored data between device consumes network and/or communication bandwidth in the system 12500, and/or operations on sensor data such as processing, compression, statistical analysis, summarization, and/or provision of alerts consumes processor cycles as well as memory to support operations such as buffer files, intermediate data, and the like. Accordingly, improved or optimal configuration and/or updating of the data storage

profile 12532 provides for lower utilization of system resources and/or allows for the storage of sensor data with higher resolution, over longer time frames, and/or from a larger number of sensors.

[1393] Referencing Figure 129, an example apparatus 12520 for self-organizing data storage for a data collector for an industrial system is depicted. An example apparatus 12520 includes a controller, such as controller 12512. The example controller includes a sensor communication circuit 12522 that interprets a number of sensor data values 12542, and a sensor data storage profile circuit 12524 that determines a data storage profile 12532. The data storage profile 12532 includes a data storage plan for the number of sensor data values 12542. The data storage plan includes how much of the sensor data values 12542 is stored initially (e.g., as the data is sampled, and/or after initial transmission to a data controller 12508, plant computer 12510, controller 12512, and/or cloud-computing device 12514). The example data storage profile 12532 includes a plan for the transmission of data, which may be according to a time, a process stage, operating conditions of the system 12500 and/or a network related to the system, as well as the communication conditions of devices within the system 12500.

[1394] For example, data from a temperature sensor may be planned to be stored locally on a sensor having storage capacity, and transmitted in bursts to a data controller. The data controller may be instructed to transmit the sensor data to the cloud computing device on a schedule, for example as the data controller memory reaches a threshold, as network communication capacity is available, at the conclusion of a process, and/or upon request. Additionally or alternatively, data from the sensors may be changed on a device or upon transfer of the data (e.g., just before transfer, just after transfer, or on a schedule). For example, the data storage profile 12532 may describe storing high resolution, high precision, and/or high-sampling rate data, and reducing the storage of the data set after a period of time, a selected event, and/or confirmation of a successful process or that the high resolution data is no longer needed. Accordingly, higher resolution data and/or data from a large number of sensors may be available for utilization, such as by a sensor fusion operation or the like, while the long-term memory utilization is also managed. Each of the sensor data sets may be treated individually for memory storage characteristics, and/or sensors may be grouped for similar treatment (e.g., sensors having similar data characteristics and/or impact on the system, sensors cooperating in a sensor fusion operation, a group of sensors utilized for a model or a virtual sensor, etc.). In certain example embodiments, sensor data from a single sensor may be treated distinctly according to an update of the data storage profile 12532, a time or process stage at which the data is taken, and/or a system condition such as a network issue, a fault condition, or the like. Additionally or alternatively, a single set of sensor data may be stored in multiple places in the system, for example where the same data is utilized in several separate sensor fusion operations, and the resource consumption from storing multiple sets of the same data is lower than a processor or network utilization to utilize a single stored data set in several separate processes.

[1395] Referencing Figure 133, various aspects of an example data storage profile 12532 are depicted. The example data storage profile 12532 includes aspects of the data storage profile 12532 that may be included as additional or alternative aspects of the data storage profile 12532 relative to the storage location definition 12534, the storage time definition, and/or the storage time

definition 12536, data resolution description 12540, and/or may be included as aspects of these. Any one or more of the factors or parameters relating to storage depicted in Figure 133 may be included in a data storage profile 12532 and/or managed by a self-organizing storage system (e.g., system 12500 and/or controller 12532). The self-organizing storage system may manage or optimize any such parameters or factors noted throughout this disclosure, individually or in combination, using an expert system, which may involve a rule-based optimization, optimization based on a model of performance, and/or optimization using machine learning/artificial intelligence, optionally including deep learning approaches, or a hybrid or combination of the above. In example embodiments, an example data storage profile 12532 includes a storage type plan 12576 or profile that accounts for or specifies a type of storage, such as based on the underlying physical media type of the storage, the type of device or system on which storage resides, the mechanism by which storage can be accessed for reading or writing data, or the like. For example, a storage media plan 12578 may specify or account for use of tape media, hard disk drive media, flash memory media, non-volatile memory, optical media, one-time programmable memory, or the like. The storage media plan may account for or specify parameters relating to the media, including capabilities such as storage duration, power usage, reliability, redundancy, thermal performance factors, robustness to environmental conditions (such as radiation or extreme temperatures), input/output speeds and capabilities, writing speeds, reading speeds, and the like, or other media specific parameters such as data file organization, operating system, read-write life cycle, data error rates, and/or data compression aspects related to or inherent to the media or media controller. A storage access plan 12580 or profile may specify or account for the nature of the interface to available storage, such as database storage (including relational, object-oriented, and other databases, as well as distributed databases, virtual machines, cloud-based databases, and the like), cloud storage (such as S3™ buckets and other simple storage formats), stream-based storage, cache storage, edge storage (e.g., in edge-based network nodes), on-device storage, server-based storage, network-attached storage or the like. The storage access plan or profile may specify or account for factors such as the cost of different storage types, input/output performance, reliability, complexity, size, and other factors. A storage protocol plan 12582 or profile may specify or account for a protocol by which data will be transmitted or written, such as a streaming protocol, an IP-based protocol, a non-volatile memory express protocol, a SATA protocol or other network-attached storage protocol, a disk-attached storage protocol, an Ethernet protocol, a peered storage protocol, a distributed ledger protocol, a packet-based storage protocol, a batch-based storage protocol, a metadata storage protocol, a compressed storage protocol (using various compression types, such as for packet-based media, streaming media, lossy or lossless compression types, and the like), or others. The storage protocol plan may account for or specify factors relating to the storage protocol, such as input/output performance, compatibility with available network resources, cost, complexity, data processing required to implement the protocol, network utilization to support the protocol, robustness of the protocol to support system noise (e.g., EM, competing network traffic, interruption frequency of network availability), memory utilization to implement the protocol (such as: as-stored memory utilization, and/or intermediate memory

utilization in creating or transferring the data), and the like. A storage writing protocol 12584 plan or profile may specify or account for how data will be written to storage, such as in file form, in streaming form, in batch form, in discrete chunks, to partitions, in stripes or bands across different storage locations, in streams, in packets or the like. The storage writing protocol may account for or specify parameters and factors relating to writing, such as input speed, reliability, redundancy, security, and the like. A storage security plan 12586 or profile may account for or specify how storage will be secured, such as availability or type of password protection, authentication, permissioning, rights management, encryption (of the data, of the storage media, and/or of network traffic on the system), physical isolation, network isolation, geographic placement, and the like. A storage location plan 12588 or profile may account for or specify a location for storage, such as a geolocation, a network location (e.g., at the edge, on a given server, or within a given cloud platform or platforms), or a location on a device, such as a location on a data collector, a location on a handheld device (such as a smart phone, tablet, or personal computer of an operator within an environment), a location within or across a group of devices (such as a mesh, a peer-to-peer group, a ring, a hub-and-spoke group, a set of parallel devices, a swarm of devices (such as a swarm of collectors), or the like), a location in an industrial environment (such as or within an storage element of an instrumentation system of or for a machine, a location on an information technology system for the environment, or the like), or a dedicated storage system, such as a disk, dongle, USB device, or the like. A storage backup plan 12590 or profile may account for or specify a plan for backup or redundancy of stored data, such as indicating redundant locations and managing any or all of the above factors for a backup storage location. In certain example embodiments, the storage security plan 12586 and/or storage backup plan 12590 may specify parameters such as data retention, long-term storage plans (e.g., migrate the stored data to a different storage media after a period of time and/or after certain operations in the system are performed on the data), physical risk management of the data and/or storage media (e.g., provision of the data in multiple geographic regions having distinct physical risk parameters, movement of the data when a storage location experiences a physical risk, refreshing the data according to a predicted life cycle of a long-term storage media, etc.).

[1396] The example controller 12512 further includes a sensor data storage implementation circuit 12526 that stores at least a portion of the number of sensor data values in response to the data storage profile 12532. An example controller 12512 includes the data storage profile 12532 having a storage location definition 12534 corresponding to at least one of the number of sensor data values 12542, including at least one location such as: a sensor storage location (e.g., data stored for a period of time on the sensor, and/or on a portable device for a user 12518 in proximity to the industrial system 12502 where the portable device is adapted by the system as a sensor), a sensor communication device storage location (e.g., a data controller 12508, MUX device, smart sensor in communication with other sensors, and/or on a portable device for a user 12518 in proximity to the industrial system 12502 or a network of the industrial system 12502 where the portable device is adapted by the system as a communication device to transfer sensor data between components in the system, etc.), a regional network storage location (e.g., on a plant computer 12510 and/or

controller 12512), and/or a global network storage location (e.g., on a cloud computing device 12514).

[1397] An example controller 12512 includes the data storage profile 12532 including a storage time definition 12536 corresponding to at least one of the number of sensor data values 12542, including at least one time value such as: a time domain description over which the corresponding at least one of the number of sensor data values is to be stored (e.g., times and locations for the data, which may include relative time to some aspect such as the time of data sampling, a process stage start or stop time, etc., or an absolute time such as midnight, Saturday, the first of the month, etc.); a time domain storage trajectory including a number of time values corresponding to a number of storage locations over which the corresponding at least one of the number of sensor data values is to be stored (e.g., the flow of the sensor data through the system across a number of devices, with the time for each storage transfer including a relative or absolute time description); a process description value over which the corresponding at least one of the number of sensor data values is to be stored (e.g., including a process description and the planned storage location for data values during the described process portion; the process description can include stages of a process, and identification of which process is related to the storage plan, and the like); and/or a process description trajectory including a number of process stages corresponding to a number of storage locations over which the corresponding at least one of the number of sensor data values is to be stored (e.g., the flow of the sensor data through the system across a number of devices, with process stage and/or process identification for each storage transfer).

[1398] An example controller 12512 includes the data storage profile 12532 including a data resolution description 12540 corresponding to at least one of the number of sensor data values 12544, where the data resolution description 12540 includes a value such as: a detection density value corresponding to the at least one of the number of sensor data values (e.g., detection density may be time sampling resolution, spatial sampling resolution, precision of the sampled data, and/or a processing operation to be applied that may affect the available resolution, such as filtering and/or lossy compression of the data); a detection density value corresponding to a more than one of the number of the sensor data values (e.g., a group of sensors having similar detection density values, a secondary data value determined from a group of sensors having a specified detection density value, etc.); a detection density trajectory including a number of detection density values of the at least one of the number of sensor data values, each of the number of detection density values corresponding to a time value (e.g., any of the detection density concepts combined with any of the time domain concepts); a detection density trajectory including a number of detection density values of the at least one of the number of sensor data values, each of the number of detection density values corresponding to a process stage value (e.g., any of the detection density concepts combined with any of the process description or stage concepts); and/or a detection density trajectory comprising a number of detection density values of the at least one of the number of sensor data values, each of the number of detection density values corresponding to a storage location value (e.g., detection density can be varied according to the device storing the data).

[1399] An example sensor data storage profile circuit 12524 further updates the data storage profile 12532 after the operations of the sensor data storage implementation circuit 12526, where the sensor data storage implementation circuit 12526 further stores the portion of the number of sensor data values 12544 in response to the updated data storage profile 12532. For example, during operations of a system at a first point in time, the sensor data storage implementation circuit 12526 utilizes a currently existing data storage profile sensor data storage implementation circuit 12526, which may be based on initial estimates of the system performance, desired data from an operator of the system, and/or from a previous operation of the sensor data storage profile circuit 12524. During operations of the system, the sensor data storage implementation circuit 12526 stores data according to the data storage profile 12532, and the sensor data storage profile circuit 12524 determines parameters for the data storage profile 12532 which may result in improved performance of the system. An example sensor data storage profile circuit 12524 tests various parameters for the data storage profile 12532, for example utilizing a machine learning optimization routine, and upon determining that an improved data storage profile 12532 is available, the sensor data storage profile circuit 12524 provides the updated data storage profile 12532 which is utilized by the sensor data storage implementation circuit 12526. In certain example embodiments, the sensor data storage profile circuit 12524 may perform various operations such as supplying an intermediate data storage profile 12532 which is utilized by the sensor data storage implementation circuit 12526 to produce real-world results, applies modeling to the system (either first principles modeling based on system characteristics, a model utilizing actual operating data for the system, a model utilizing actual operating data for an offset system, and/or combinations of these) to determine what an outcome of a given data storage profile 12532 will be or would have been (including, for example, taking extra sensor data beyond what is utilized to support a process operated by the system), and/or applying randomized changes to the data storage profile 12532 to ensure that an optimization routine does not settle into a local optimum or non-optimal condition.

[1400] An example sensor data storage profile circuit 12524 further updates the data storage profile 12532 in response to external data 12544 and/or cloud-based data 12538, including data such as: an enhanced data request value (e.g., an operator, model, optimization routine, and/or other process requests enhanced data resolution for one or more parameters); a process success value (e.g., indicating that current storage practice provides for sufficient data availability and/or system performance; and/or that current storage practice may be over-capable, and one or more changes to reduce system utilization may be available); a process failure value (e.g., indicating that current storage practices may not provide for sufficient data availability and/or system performance, which may include additional operations or alerts to an operator to determine whether the data transmission and/or availability contributed to the process failure); a component service value (e.g., an operation to adjust the data storage to ensure higher resolution data is available to improve a learning algorithm predicting future service events, and/or to determine which factors may have contributed to premature service); a component maintenance value (e.g., an operation to adjust the data storage to ensure higher resolution data is available to improve a learning algorithm predicting future maintenance events, and/or to determine which factors may have contributed to premature

maintenance); a network description value (e.g., a change in the network, for example by identification of devices, determination of protocols, and/or as entered by a user or operator, where the network change results in a capability change and potentially a distinct optimal storage plan for sensor data); a process feedback value (e.g., one or more process conditions detected); a network feedback value (e.g., one or more network changes as determined by actual operations of the network – e.g., a loss or reduction in communication of one or more devices, a network communication volume change, a transmission noise value change on the network, etc.); a sensor feedback value (e.g., metadata such as a sensor fault, capability change; and/or based on the detected data from the system, for example an anomalous reading, rate of change, or off-nominal condition indicating that enhanced or reduced resolution, sampling time, etc. should change the storage plan); and/or a second data storage profile, where the second data storage profile was generated for an offset system.

[1401] An example storage planning circuit 12528 determines a data configuration plan 12546 and updates the data storage profile 12532 in response to the data configuration plan 12546, where the sensor data storage implementation circuit 12526 further stores at least a portion of the number of sensor data values in response to the updated data storage profile 12532. An example data configuration plan 12546 includes a value such as: a data storage structure value (e.g., a data type, such as integer, string, a comma delimited file, how many bits are committed to the values, etc.); a data compression value (e.g., whether to compress data, a compression model to use, and/or whether segments of data can be replaced with summary information, polynomial or other curve fit summarizations, etc.); a data write strategy value (e.g., whether to store values in a distributed manner or on a single device, which network communication and/or operating system protocols to utilize); a data hierarchy value (e.g., which data is favored over other data where storage constraints and/or communication constraints will limit the stored data – the limits may be temporal, such as data will not be in the intended location at the intended time, or permanent, such as some data will need to be compressed in a lossy manner, and/or lost); an enhanced access value determined for the data (e.g., the data is of a type for reports, searching, modeling access, and/or otherwise tagged, where enhanced access includes where the data is stored for scope of availability, indexing of data, summarization of data, topical reports of data, which may be stored in addition to the raw or processed sensor data); and/or an instruction value corresponding to the data (e.g., a placeholder indicating where data can be located, an interface to access the data, metadata indicating units, precision, time frames, processes in operation, faults present, outcomes, etc.).

[1402] It can be seen that the provision of control over data flow and storage through the system allows for improvement generally, and movement toward optimization over time, of data management throughout the system. Accordingly, more data of a higher resolution can be accumulated, and in a more readily accessible manner, than previously known systems with fixed or manually configurable data storage and flow for a given utilization of resources such as storage space, communication bandwidth, power consumption, and/or processor execution cycles. Additionally, the system can respond to process variations that affect the optimal or beneficial parameters for controlling data flow and storage. One of skill in the art, having the benefit of the

disclosures herein, will recognize that combinations of control of data storage schemes with data type control and knowledge about process operations for a system create powerful combinations in certain contemplated embodiments. For example, data of a higher resolution can be maintained for a longer period and made available if a need for the data arises, without incurring the full cost of storing the data permanently and/or communicating the data throughout every layer of the system.

[1403] In example embodiments, in an underground mining inspection system, certain detailed data regarding toxic gas concentrations, temperatures, noise, etc. may need to be captured and stored for regulatory purposes, but for ongoing operational purposes, perhaps only a single data point regarding one or more toxic gases is needed periodically. In this embodiment, the data storage profile for the system may indicate that only certain sensor data aligned with regulatory needs be stored in a certain manner that is long term and optionally only available as needed, while other sensor data required operationally be stored in a more accessible manner.

[1404] In another embodiment involving automotive brakes for fleet vehicles, data regarding brake use and performance may be acquired at high resolution and stored in a first data storage that is not transmitted throughout the network, while lower resolution data are transmitted periodically and/or in near real time to a fleet control and maintenance application. Should the application or other user require higher resolution data, it may be accessed from the first data storage.

[1405] In a further embodiment of manufacturing body and frame components of trucks and cars, certain detailed data regarding paint color, surface curvature, and other quality control measures may be captured and stored at high resolution, but for ongoing operational purposes, only low resolution data regarding throughput are transmitted. In this embodiment, the data storage profile for the system may indicate that only certain sensor data aligned with quality control needs be stored in a certain manner that is long term and optionally only available as needed, while other sensor data required operationally be stored in a more accessible manner.

[1406] In another example, data types, resolution, and the like can be configured and changed as the data flows through the system, according to values that are beneficial for the individual components handling the data, according to the utilized networking resources for the data, and/or according to accompanying data (e.g., a model, virtual sensor, and/or sensor fusion operation) where higher capability data would not improve the precision of the process utilizing the accompanying data.

[1407] In example embodiments, in rail condition monitoring systems, as rail condition data are acquired, each component of the system may require different resolutions of the same data. Continuing with this example, as real-time rail traffic data are acquired, these data may be stored and/or transmitted at low resolution in order to quickly disseminate the data throughout the system, while utilization and load data may be stored and utilized at higher resolution to track rail use fees and need for rail maintenance at a more granular level.

[1408] In another embodiment of a hydraulic pump operating in a tractor, as the tractor is in the field and does not have access to a network, data from on-board sensors may be acquired and stored

in a local manner on the tractor at low resolution, but when the tractor regains access, data may be acquired and transmitted at high resolution.

[1409] In yet another embodiment of an actuator in a robotic handling unit in an automotive plant, data regarding the actuator may flow into multiple downstream systems, such as a production tracking system that utilizes the actuator data alone and an energy efficiency tracking system that utilizes the data in a sensor fusion with data from environmental sensors. Resolution of the actuator data may be configured differently as it is transmitted to each of these systems for their disparate uses.

[1410] In still another embodiment of a generator in a mine, data may be acquired regarding the performance of the generator, carbon monoxide levels near the generator and a cost for running the generator. Each component of a control system overseeing the mine may require different resolutions of the same data. Continuing with this example, as carbon monoxide data are acquired, these data may be stored and/or transmitted at low resolution in order to quickly disseminate the data throughout the system in order to properly alert workers. Performance and cost data may be stored and utilized at higher resolution to track economic efficiency and lifetime maintenance needs.

[1411] In an additional embodiment, sensors on a truck's wheel end may monitor lubrication, noise (e.g., grinding, vibration) and temperature. While in the field, sensor data may be transmitted remotely at low resolution for remote monitoring, but when within a threshold distance from a fleet maintenance facility, data may be transmitted at high resolution.

[1412] In another example, accompanying information for the data allows for efficient downstream processing (e.g., by a downstream device or process accessing the data) including unpackaging the data, readily determining where related higher capability data may be present in the system, and/or streamlining operations utilizing the data (e.g., reporting, modeling, alerting, and/or performing a sensor fusion or other system analysis). An embodiment includes storing high capability (e.g., high-sampling rate, high precision, indexed, etc.) in a first storage device in the system (e.g., close to the sensors in the network layer to preserve network communication resources) and sending lower capability data up the network layers (e.g., to a cloud-computing device), where the lower capability data includes accompanying information to access the stored high capability data, including accompanying data that may be accessible to a user (e.g., a header, message box, or other organically interfaceable accompanying data) and/or accessible to an automated process (e.g., structured data, XML, populated fields, or the like) where the process can utilize the accompanying data to automatically request, retrieve, or access the high capability data. In certain example embodiments, accompanying data may further include information about the content, precision, sampling time, calibrations (e.g., de-bouncing, filtering, or other processing applied) such that an accessing component or user can determine without retrieving the high capability data whether such data will meet the desired parameters.

[1413] In example embodiments, vibration noise from vibration sensors attached to vibrators on an assembly line may be stored locally in a high resolution format while a low resolution version of the same data with accompanying information regarding the availability of ambient and local

noise data for a sensor fusion may be transmitted to a cloud-based server. If a resident process on the server requires the high resolution data, such as a machine learning process, the server may retrieve the data at that time.

[1414] In another embodiment of an airplane engine, performance data aggregated from a plurality of sensors may be transmitted while in flight along with accompanying information to a remote site. The accompanying information, such as a header with metadata relating to historical plane information, may allow the remote site to efficiently analyze the performance data in the context of the historical data without having to access additional databases.

[1415] In a further embodiment of a coal crusher in a power generation facility, data accompanying low quality sensor data regarding the size of coal exiting the crusher may include information about the precision in the size measurement such that a technician can determine if the higher resolution data are needed to confirm a determination that the crusher needs to come offline for maintenance.

[1416] In yet a further embodiment of a drilling machine or production platform employed in oil and gas production, high capability data may be acquired and stored locally regarding parameters of the drill's and platform's operation, but only low capability data are transmitted off-site to conserve bandwidth. Along with the low capability data, accompanying information may include instructions on how an automated off-site process can automatically access the high capability data in the event that it is required.

[1417] In still a further embodiment, temperature sensors on a pump employed in oil & gas production or mining may be stored locally in a high resolution format while a low resolution version of the same data with accompanying information regarding the availability of noise and energy use data for a sensor fusion may be transmitted to a cloud-based server. If a resident process on the server requires the high resolution data, such as a machine learning process, the server may retrieve the data at that time.

[1418] In another embodiment of a gearbox in an automatic robotic handling unit or an agricultural setting, performance data aggregated from a plurality of sensors may be transmitted while in use along with accompanying information to a remote site. The accompanying information, such as a header with metadata relating to historical gearbox information, may allow the remote site to efficiently analyze the performance data in the context of the historical data without having to access additional databases.

[1419] In a further embodiment of a ventilation system in a mine, data accompanying low quality sensor data regarding the size of particulates in the air may include information about the precision in the size measurement such that a technician can determine if the higher resolution data are needed to confirm a determination that the ventilation system requires maintenance.

[1420] In yet a further embodiment of a rolling bearing employed in agriculture, high capability data may be acquired and stored locally regarding parameters of the rolling bearing's operation, but only low capability data are transmitted off-site to conserve bandwidth. Along with the low capability data, accompanying information may include instructions on how an automated off-site process can automatically access the high capability data in the event that it is required.

[1421] In a further embodiment of a stamp mill in a mine, data accompanying low quality sensor data regarding the size of mineral deposits exiting the stamp mill may include information about the precision in the size measurement such that a technician can determine if the higher resolution data are needed to confirm a determination that the stamp mill requires a change in an operation parameter.

[1422] Referencing Figure 130, an example storage time definition 12536 is depicted. The example storage time definition 12536 depicts a number of storage locations 12556 corresponding to a number of time values 12558. It is understood that any values such as storage types, storage media, storage access, storage protocols, storage writing values, storage security, and/or storage backup values, may be included in the storage time definition 12536. Additionally or alternatively, an example storage time definition 12536 may include process operations, events, and/or other values in addition to or as an alternative to time values 12558. The example storage time definition 12536 depicts movement of related sensor data to a first storage location 12550 over a first time interval, to a second storage location 12552 over a second time internal, and to a third storage location 12554 over a third time interval. The storage location values 12550, 12552, 12554 are depicted as an integral selection corresponding to planned storage locations, but additionally or alternatively the values may be continuous or discrete, but not necessarily integral values. For example, a storage location value 12550 of "1" may be associated with a first storage location, and a storage location value 12550 of "2" may be associated with a second storage location, where a value between "1" and "2" has an understood meaning – such as a prioritization to move the data (e.g., a "1.1" indicates that the data should be moved from "2" to "1" with a relatively high priority compared to a "1.4"), a percentage of the data to be moved (e.g., to control network utilization, memory utilization, or the like during a transfer operation), and/or a preference for a storage location with alternative options (e.g., to allow for directing storage location, and inclusion in a cost function such that storage location can be balanced with other constraints in the system). Additionally or alternatively, the storage time definition 12536 can include additional dimensions (e.g., changing protocols, media, security plans, etc.) and/or can include multiple options for the storage plan (e.g., providing a weighted value between 2, 3, 4, or more storage locations, protocols, media, etc. in a triangulated or multiple-dimension definition space).

[1423] Referencing Figure 131 an example data resolution description 12540 is depicted. The example data resolution description 12540 depicts a number of data resolution values 12562 corresponding to a number of time values 12564. It is understood that any values such as storage types, storage media, storage access, storage protocols, storage writing values, storage security, and/or storage backup values, may be included in the data resolution description 12540. Additionally or alternatively, an example data resolution description 12540 may include process operations, events, and/or other values in addition to or as an alternative to time values 12558. The example data resolution description 12540 depicts changes in the resolution of stored related sensor data resolution values 12560 over time intervals, for example operating at a low resolution initially, stepping up to a higher resolution (e.g., corresponding to a process start time), to a high resolution value (e.g., during a process time where the process is significantly improved by high resolution

of the related sensor data), and to a low resolution value (e.g., after a completion of the process). The example depicts a higher resolution before the process starts than after the process ends as an illustrative example, but the data resolution description 12540 may include any data resolution trajectory. The data resolution values 12560 are depicted as integral selections corresponding to planned data resolutions, but additionally or alternatively the values may be continuous or discrete, but not necessarily integral values. For example, data resolution values 12560 of "1" may be associated with a first data resolution (e.g., a specific sampling time, byte resolution, etc.), and a data resolution values 12560 of "2" may be associated with a second data resolution, where a value between "1" and "2" has an understood meaning – such as a prioritization to sample at the defined resolution (e.g., a "1.1" indicates the data should be taken at a sampling rate corresponding to "1" with a relatively high priority compared to a "1.3", and/or at a sampling rate 10% of the way between the rate between "1" and "2"), and/or a preference for a data resolution with alternative options (e.g., to allow for sensor or network limitations, available sensor communication devices such as a data controller, smart sensor, or portable device taking the data from the sensor, and/or inclusion in a cost function such that data resolution can be balanced with other constraints in the system). Additionally or alternatively, the data resolution description 12540 can include additional dimensions (e.g., changing protocols, media, security plans, etc.) and/or can include multiple options for the data resolution plan (e.g., providing a weighted value between 2, 3, 4, or more data resolution values, protocols, media, etc. in a triangulated or multiple-dimension definition space).

[1424] An example system 12500 further includes a haptic feedback circuit 12530 that determines a haptic feedback instruction 12548 in response to at least one of the number of sensor values 12542 and/or the data storage profile 12532, and a haptic feedback device 12516 responsive to the haptic feedback instruction 12548. Example and non-limiting haptic feedback instructions 12548 include an instruction such as: a vibration command; a temperature command; a sound command; an electrical command; and/or a light command. Example and non-limiting operations of the haptic feedback circuit 12530 include feedback that data is stored or being stored on the haptic feedback device 12516 and/or on a portable device associated with the user 12518 in communication with the haptic feedback device 12516 (e.g., user 12518 traverses through the system 12500 with a smart phone, which the system 12500 utilizes to store sensor data, and provides a haptic feedback instructions 12548 to notify the user 12518 that the smart phone is currently being utilized by the system 12500, for example allowing the user 12518 to remain in communication with the sensor, data controller, or other transmitting device, and/or allowing the user to actively cancel or enable the data transfer). Additionally or alternatively, the haptic feedback device 12516 may be the smart phone (e.g., utilizing vibration, sound, light, or other haptic aspects of the smart phone), and/or the haptic feedback device 12516 may include data storage and/or communication capabilities.

[1425] In certain example embodiments, the haptic feedback circuit 12530 provides a haptic feedback instruction 12548 as an alert or notification to the user 12518, for example to alert or notify the user 12518 that a process has commenced or is about to start, that an off-nominal operation is detected or predicted, that a component of the system requires or is predicted to require maintenance, that an aspect of the system is in a condition that the user 12518 may want to be

aware of (e.g., a component is still powered, has high potential energy of any type, is at a high pressure, and/or is at a high temperature where the user 12518 may be in proximity to the component), that a data storage related aspect of the system is in a noteworthy condition (e.g., a data storage component of the system is at capacity, out of communication, is in a fault condition, has lost contact with a sensor, etc.), to request a response from the user 12518 (e.g., an approval to start a process, data transfer, process rate change, clear a fault, etc.) In certain example embodiments, the haptic feedback circuit 12530 configures the haptic feedback instruction 12548 to provide an intuitive feedback to the user 12518. For example, an alert value may provide a more rapid, urgent, and/or intermittent vibration mode relative to an informational notification; a temperature based alert or notification may utilize a temperature based haptic feedback (e.g., an overtemperature vessel notification may provide a warm or cold haptic feedback) and/or flashing a color that is associated with the temperature (e.g., flashing red for an overtemperature or blue for an under-temperature); an electrically based notification may provide an electrically associated haptic feedback (e.g., a sound associated with electricity such as a buzzing or sparking sound, or even a mild electrical feedback such as when a user is opening a panel for a component that is still powered); providing a vibration feedback for a bearing, motor, or other rotating or vibrating component that is operating off-nominally; and/or providing a requested feedback to the user based upon sensed data (e.g., transmitting a vibration profile to the haptic feedback device that is analogous to the detected vibration in a requested component for example allowing an expert user to diagnose the component without physical contact; providing a haptic feedback for a requested component for example if the user is double checking a lockout/tagout operation before entering a component, opening a panel, and/or entering a potentially hazardous area). The provided examples for operations of the haptic feedback circuit 12530 are non-limiting illustrations.

[1426] Referencing Figure 132, an example apparatus for data collection in an industrial environment 12566 includes a controller 12512 a sensor communication circuit 12522 that interprets a number of sensor data values 12542, a sensor data storage profile circuit 12524 that determines a data storage profile 12532, where the data storage profile 12532 includes a data storage plan for the number of sensor data values 12542, and a network coding circuit 12568 that provides a network coding value 12570 in response to the number of sensor data values 12542 and the data storage profile 12532. The controller 12512 further includes a sensor data storage implementation circuit 12526 that stores at least a portion of the number of sensor data values 12542 in response to the data storage profile 12532 and the network coding value 12570. The network coding value 12570 includes, without limitation, network encoding for data transmission, such as packet sizing, distribution, combinations of sensor data within packets, encoding and decoding algorithms for network data and communications, and/or any other aspects of controlling network communications throughout the system. In certain example embodiments, the network coding value 12570 includes a linear network coding algorithm, a random linear network coding algorithm, and/or a convolutional code. Additionally or alternatively, the network coding circuit 12568 provides scheduling and/or synchronization for network communication devices of the system, and can include separate scheduling and/or synchronization for separate networks in the

system. The network coding circuit 12568 schedules the network coding value 12570 throughout the system according to the data volumes, transfer rates, and network utilization, and alternatively or additionally performs a self-learning and/or machine learning operation to improve or optimize network coding. For example, a sensor having a single low-volume data transfer to a data controller may utilize TCP/IP packet communication to the data controller without linear network coding, while higher volume aggregated data transfer from the data controller to another system component (e.g., the controller 12532) may utilize linear network coding. The example network coding circuit 12568 adjusts the network coding value 12570 in real time for the components in the system to optimize or improve transfer rates, power utilization, errors and lost packets, and/or any other desired parameters. For example, a given component may have resulting low transfer rates but a large available memory, while a downstream component has a lower available memory (potentially relative to the data storage expectation for that component), and accordingly a complex network coding value 12570 for the given component may not result in improved throughput of data throughout the system, while a network coding value 12570 enhancing throughput for the downstream component may justify the processing overhead for a more complex network coding value 12570.

[1427] An example system includes the network coding circuit 12568 further determining a network definition value 12572, and providing the network coding value 12570 further in response to the network definition value 12572. Example network definition values 12572 include values such as: a network feedback value (e.g., transfer rates, up time, synchronization availability, etc.); a network condition value (e.g., presence of noise, transmission/receiver capability, drop-outs, etc.); a network topology value (e.g., the communication flow and connectivity of devices; operating systems, protocols, and storage types of devices; available computing resources on devices; the location and function of devices in the system); an intermittently available network device value (e.g., a known or observed availability for the device over time or process stage; predicted availability of the device; prediction of known noise factors for the device, such as process operations that reduce device availability); and/or a network cost description value (e.g., resource utilization of the device, including relative cost or impact of processing, memory, and/or communication resources; power utilization and cost of power consumption for devices; available power for the device and a cost description for externalities related to consuming the power - such as for a battery where the power itself may not be expensive but the power in the specific location has a cost associated with replacement, including availability or access to the device during operations).

[1428] An example system includes the network coding circuit 12568 further providing the network coding value 12570 such that the sensor data storage implementation circuit stores a first portion of the number of sensor data values 12542 utilizing a first network coding value 12570, and a second portion of the number of sensor data values 12542 utilizing a second network coding value 12570 (e.g., the network coding values 12570 can vary with the data being transmitted, the transmitting device, and/or over time or process stage). Example and non-limiting network coding values include: a network type selection (e.g., public, private, wireless, wired, intranet, external,

internet, cellular, etc.), a network selection(e.g., which one or more of an available number of networks will be utilized), a network coding selection (e.g., packet definitions, encoding techniques, linear, randomized linear, convolution, triangulated, etc.), a network timing selection (e.g., synchronization and sequencing of data transmissions between devices), a network feature selection (e.g., turning on or off network support devices or repeaters; enabling, disabling, or adjusting security selections; increasing or decreasing a power of a device, etc.), a network protocol selection (e.g., TCP/IP, FTP, Wi-Fi, BLUETOOTH™, Ethernet, and/or routing protocols); a packet size selection (including header and/or parity information); and/or a packet ordering selection (e.g., determining how to transmit the various sensor information that may be on a device, and/or determining the packet to data value correspondence). An example network coding circuit 12568 further adjusts the network coding value 12570 to provide an intermediate network coding value (e.g., as a test coding value on the system, and/or as a modeled coding value being run off-line), to compare a performance indicator 12574 corresponding to each of the network coding value 12570 and the intermediate network coding value, and to provide an updated network coding value (e.g., as the network coding value 12570) in response to the comparison of the performance indicators 12574.

[1429] An example system includes an industrial system having a number of components, and a number of sensors each operatively coupled to at least one of the number of components. The number of sensors provide a number of sensor values, and the system further includes a number of organizing structures such as a controller, a data collector, a plant computer, a cloud-based server and/or global computing device, and/or a network layer, where the organizing structures are configured for self-organizing storage of at least a portion of the number of sensor values. For example, operations of the controller 12512 provide for storage and distribution of sensor data values to reduce consumption of resources (processor, network, and/or memory) for storing sensor data. The self-organizing operations include management of the stored sensor data over time, including providing sensor information to system components in time to complete operations therefore (e.g., control, improvement, modeling, and/or machine learning for process operations of the system). Additionally, data security, including long-term security due to storage media, geographic, and/or unauthorized access, is considered throughout the data storage life cycle. An example system further includes the organizing structures providing enhanced resolution of the number of sensor values in response to at least one of an enhanced data request value or an alert value corresponding to the industrial system. The system provides enhanced resolution by controlling the storage processes to address system impact, including keeping lower resolution, summary, or other accessibility data available, and storing higher resolution data in a lower resource utilization manner which is available upon request and/or at a time appropriate to system operations. Example enhanced resolution includes: an enhanced spatial resolution, an enhanced time domain resolution, a greater number of the number of sensor values than a standard resolution of the number of sensor values, and/or a greater precision of at least one of the number of sensor values than a standard resolution of the number of sensor values. An example system further includes a network layer, where the organizing structures are configured for self-organizing

network coding for communication of the number of sensor values on the network layer. An example system further includes a haptic feedback device of a user in proximity to at least one of the industrial system or the network layer, and where the organizing structures are configured for providing haptic feedback to the haptic feedback device, and/or for configuring the haptic feedback to provide an intuitive alert to the user.

[1430] In example embodiments, the expert system uses at least one of a rule and a model to set a parameter of the error correction.

[1431] In example embodiments, the expert system is a machine learning system that iteratively configures at least one of a set of inputs, a set of weights, and a set of functions based on feedback relating to at least one of the data paths.

[1432] As depicted in Figure 134, a cloud platform for supporting deployments of devices in the Internet of Things (IoT), such as within industrial environments, may include various components, modules, services, elements, applications, interfaces, and other elements (collectively referred to as the "cloud platform 13000"), which may include a policy automation engine 13002 and a data marketplace 13008. The cloud platform 13000 may include, integrate with, or connect to various devices 13006, a cloud computing environment 13068, data pools 13070, data collectors 13020 and sensors 13024. The cloud platform 13000 may also include systems and capabilities for self-organization 13012, machine learning 13014 and rights management 13016.

[1433] Within the cloud platform 13000, various components may be deployed in a wide range of architectures and arrangements. In example embodiments, devices 13006 may connect to, integrate with, or be deployed within a cloud computing environment 13068, the policy automation engine 13002, the data marketplace 13008, the data collectors 13020, as well as systems and capabilities for self-organization 13012, machine learning 13014 and rights management 13016. Devices 13006 may connect to or integrate with the policy automation engine 13002, data marketplace 13008, data collectors 13020 and systems or capabilities for self-organization 13012, machine learning 13014 and rights management 13016, either directly or through the cloud computing environment 13068.

[1434] Devices 13006 may be IoT devices, including IoT devices, such as for collecting, exchanging and managing information relating to machines, personnel, equipment, infrastructure elements, components, parts, inventory, assets, and other features of a wide range of industrial environments, such as those described throughout this disclosure. Devices 13006 may also connect via various protocols 13004, such as networking protocols, streaming protocols, file transfer protocols, data transformation protocols, software operating system protocols, and the like. Devices may connect to the policy automation engine 13002, such as for executing policies that may be deployed within the cloud platform 13000, such as governing activities, permissions, rules, and the like within the platform 13000. Devices 13006 may also connect to data streams 13010 within the data marketplace 13008.

[1435] Data pools 13070 may connect to or integrate with the cloud computing environment 13068, data collectors 13020 and the data marketplace 13008, policy automation engine 13002, self-organization 13012, machine learning 13014 and rights management 13016 capabilities. Data

pools 13070 may be included within the cloud computing environment 30 or be external to the cloud computing environment 13068. As a result, connections to the data pools 13070 may be made directly to the data pools 13070, through cloud connections to the data pools 13070 or through a combination of direct and cloud connections to the data pools 13070. Data pools 13070 may also be included within the data marketplace 13008 or external to the data marketplace 13008.

[1436] Data pools 13070 may include a multiplexer (MUX) 13022 and also connect to self-organization 13012, machine learning 13014 and rights management capabilities. The MUX 13022 may connect to sensors 13024, collect data from sensors 13024 and integrate data collected from sensors 13024 into a single set of data. In an exemplary and non-limiting embodiment, data pools 13070, data collectors 13020 and sensors 13024 may be included within an industrial environment 13018.

[1437] A policy automation engine 13002 and data marketplace 13008 may be used in a variety of industrial environments 13018. Industrial environments 13018 may include aerospace environments, agriculture environment, assembly line environments, automotive environments, and chemical and pharmaceutical environments. Industrial environments 13018 may also include food processing environments, industrial component environments, mining environments, oil and gas environments, particularly oil and gas production environments, truck and car environments and the like.

[1438] Similarly, devices 13006 may include a variety of devices that may operate within the industrial environments or that may collect data with respect to other such devices. Among many examples, devices 13006 may include agitators, including turbine agitators, airframe control surface vibration devices, catalytic reactors and compressors. Devices 13006 may also include conveyors and lifters, disposal systems, drive trains, fans, irrigation systems and motors. Devices 13006 may also include pipelines, electric powertrains, production platforms, pumps, such as water pumps, robotic assembly systems, thermic heating systems, tracks, transmission systems and turbines. Devices 13006 may operate within a single industrial environment 13018 or multiple industrial environments 13018. For example, a pipeline device may operate within an oil and gas environment, while a catalytic reactor may operate in either an oil and gas production environment or a pharmaceutical environment.

[1439] The policy automation engine 13002 may be a cloud-based policy automation engine 13002. A policy automation engine 13002 may be used to create, deploy, and/or manage an interconnected set of policies 13030, rules 13028 and protocols 13004, such as policies relating to security, authorization, permissions, and the like. For example, policies may govern what users, applications, services, systems, devices, or the like may access an IoT device, may read data from an IoT device, may subscribe to a stream from an IoT device, may write data to an IoT device, may establish a network connection with an IoT device, may provision an IoT device, may collaborate with an IoT device, or the like.

[1440] The policy automation engine 13002 may generate and manage policies 13030. The policy generation engine may be the centralized policy management system for the cloud platform 13000.

[1441] Policies 13030 generated and managed by the policy automation engine 13002 may deploy a large number of rules 13028 to permit access to and use of different aspects of IoT devices. Policies 13030 may include IoT device creation policies 13032, IoT device deployment policies 13034, IoT device management policies 13036 and the like. The policies 13030 may be communicated to devices 13006 through protocols 13004 or directly from the policy automation engine 13002.

[1442] For example, in an exemplary and non-limiting embodiment, the policy automation engine 13002 may manage policies 13030 and create protocols 13004 that specify and enforce roles 13026 and permissions 13074 for workers, related to how the workers may use data provided by IoT devices. Workers may be human workers or machine workers.

[1443] In additional exemplary and non-limiting embodiments, policies 13030 may be used to automate remediation processes. Remediation processes may be performed when a system is partially disabled, when equipment fails and when an entire system may be disabled. Remediation processes may include instructions to initiate system restarts, bypass or replace equipment, notify appropriate stakeholders of the condition and the like. The policy automation engine 13002 may also include policies 13030 that specify the roles 13026 and permissions 13074 required for users 13072 to initiate or otherwise act upon the remediation or other processes.

[1444] The policy automation engine 13002 may also specify and detect conditions. Conditions may determine when policies 13030 are distributed or otherwise acted upon. Conditions may include individual conditions, sets of conditions, independent conditions, interdependent conditions, and the like.

[1445] In an exemplary and non-limiting embodiment of an independent condition, the policy automation engine 13002 may determine that the failure of a non-critical device 13006 does not require notification of the system operator. In an exemplary and non-limiting embodiment of an interdependent set of conditions, the policy automation engine 13002 may determine that the failure of two non-critical system devices 13006 does require notification of the system operator, as the failure of two non-critical system devices 13006 may be an early indicator of a possible system-wide failure.

[1446] As depicted in Figure 135, the policy automation engine 13002 may include compliance policies 13050 and fault, configuration, accounting, provisioning, and security (FCAPS) policies 13052. Policies 13030 may connect to rules 13028, protocols 13004 and policy inputs 13048.

[1447] Policies 13030 may provide input to rules 13028 and provide information related to how roles 13026, permissions 13074 and uses 130280 are defined. Policies 13030 may receive policy inputs 13048 and incorporate policy inputs 13048 as policy parameters that are included within policies 13030. Policies 13030 may provide inputs to protocols 13004 and be included within protocols 13004 that are used to create, deploy and manage devices 13006.

[1448] Compliance policies 13050 may include data ownership policies, data analysis policies, data use policies, data format policies, data transmission policies, data security policies, data privacy policies, information sharing policies, jurisdictional policies, and the like. Data transmission policies may include cross-jurisdictional data transmission policies.

[1449] Data ownership policies may indicate policies 13030 that manage who controls data, who can use data, how the data can be used and the like. Data analysis policies may indicate what data holders can do with data that they are permitted to access, as well as determine what data they can look at and what data may be combined with other data. For example, a data holder may look at aggregated user data but not individual user data. Data use policies may indicate how data may be used and under what circumstances data may be used. Data format policies may indicate standard formats and mandated formats permitted for the handling of data. Data transmission policies, including cross-jurisdictional data transmission policies, may determine the policies 13030 that specify how inter-jurisdictional and intra-jurisdictional transmission of data may be handled. Data security policies may determine how data at rest, for example stored data, as well transmitted data is required to be secured.

[1450] Data privacy policies may determine how data may or may not be shared, for example within an organization and external to an organization. Information sharing policies may determine how data may be sold, shared and under what circumstances information can be sold and shared. Jurisdictional policies may determine who controls data, when and where the data may be controlled, for data within and transmitted across boundaries.

[1451] FCAPS policies 13052 may include fault management policies, configuration management policies, accounting management policies, provisioning management policies, and security management policies. Fault management policies may specify policies 13030 used to handle device faults. Configuration management policies may specify policies used to configure devices 13006. Accounting management policies may specify policies 13030 used for device accounting purposes, such as reporting, billing and the like. Provisioning management policies may specify policies 13030 used to provision services on devices 13006. Security management policies may specify policies 13030 used to secure devices 13006.

[1452] Policy inputs 13048 may be received from a policy input interface 13046. Policy inputs 13048 may include standards-based policy inputs 13044 and other policy inputs 13048. Standards-based policy inputs 13044 may include inputs related to standard data formats, standard rule sets and other standards-related information set by standards bodies, for example.

[1453] Other policy inputs 13048 may include a wide range of information related industry-specific policies, cross-industry policies, manufacturer-specific policies, device-specific policies 13030 and the like. Policy inputs 13048 may connect to a cloud computing environment 13068 and be provided through a policy input interface 13046. The policy input interface 13046 may collect policy inputs 13048 provided by machines or entered by human operators.

[1454] As depicted in Figure 134, a data marketplace 13008 may include data streams 13010, a data marketplace input interface, data marketplace inputs 13056, a data payment allocation engine 13038, marketplace value rating engine 13040, a data brokering engine 13042, a marketplace self-organization engine 13076 and one or more data pools 13070. The data marketplace 13008 may be included within the cloud networking environment 30 or externally connected to the cloud networking environment 13068. Data pools 13070 may also be included within the cloud

networking environment 13068 or may be externally connected to the cloud networking environment 13068.

[1455] The data marketplace 13008 may connect to data pools 13070 directly, for example if the data marketplace 13008 and data pools 13070 are located in the same physical location. The data marketplace 13008 may connect to data pools 13070 via a cloud networking environment 30, for example if the data marketplace 13008 and data pools 13070 are located in different physical locations.

[1456] The data marketplace 13008 may connect to and receive inputs. The data marketplace 13008 may receive marketplace inputs through data interfaces, for example one or more data collectors 13020. The data collectors 13020 may be multiplexing data collectors. Inputs received through the data collectors 13020 may be received as one or more than one data streams 13010 from one or more than one data collectors 13020 and integrated into additional data streams 13010 by the multiplexer (MUX) 13022.

[1457] The data streams 13010 may also include data from the data pools 60. Data marketplace inputs, data streams 13010 and data pools 13070 may include metrics and measures of success of the data marketplace 13008. The metrics and measures of success of the data marketplace 13008 may then be used by the machine learning capability 13014 to configure one or more parameters of the data marketplace 13008.

[1458] Inputs may be consortia inputs 13054. Consortia inputs 13054 may be received from consortia. Consortia may include energy consortia, healthcare consortia, manufacturing consortia, smart city consortia, transportation consortia and the like. Consortia may be pre-existing consortia or new consortia.

[1459] In an exemplary and non-limiting embodiment, new consortia may be formed as a result of the data marketplace 13008 making available particular data types and data combinations. The data brokering engine 13042 may allow consortia members to trade information. The data brokering engine 13042 may allow consortia members to trade information based on information value, as calculated by the marketplace value rating engine 13040, for example.

[1460] The data marketplace 13008 may also connect to self-organization 13012, machine learning 13014 and rights management 13016 capabilities. Rights management capabilities 13016 may include rights.

[1461] Rights may include business strategy and solution rights, liaison rights 13058, marketing rights 13078, security rights 13060, technology rights 13062, testbed rights 13064 and the like. Business strategy and solution lifecycle rights may include business strategy and planning rights, industrial internet system design rights, project management rights, solution evaluation and contractual aspects rights. Liaison rights 13058 may include standards organization rights, open-source community rights, certification and testing body rights and governmental organization rights. Marketing rights 13078 may include communication rights, energy rights, healthcare rights, marketing-security rights, retail operation rights, smart factory rights and thought leadership rights. Security rights 13060 may include driving rights that drive industry consensus, promote security best practices and accelerate the adoption of security best practices.

[1462] Technology rights 13062 may include architecture rights, connectivity rights, distributed data management and interoperability rights, industrial analytics rights, innovation rights, IT/OT rights, safety rights, vocabulary rights, use case rights and liaison rights 13058. Testbed rights 13064 may include rights to implement of specific use cases and scenarios, as well as rights to produce testable outcomes to confirm that an implementation conforms to expected results, for example. Testbed rights 13064 may also include rights to explore untested or existing technologies working together, for example interoperability testing, generate new and potentially disruptive products and services and generate requirements and priorities for standards organizations, consortia and other stakeholder groups.

[1463] The rights management capability may assign different rights to different participants in the data marketplace 13008. In an exemplary and non-limiting embodiment, manufacturers or remote maintenance organizations (RMOs). Participants may be assigned rights to information based on their equipment or proprietary methods. The data marketplace 13008 may then ensure that only the appropriate data streams 13010 are made available to the market, based on the assigned rights.

[1464] The rights management capability 13016 may manage permissions to access the data in the marketplace 13008. One or more parameters of the rights management capability 13016 may be automatically configured by the machine learning capability 13014 and may be based on a metric of success of the data marketplace 13008. The machine learning engine 13014 may also use the metric and measure of success to configure a user interface. The user interface may present a data element of the user of the data marketplace 13008. The user interface may also present one or more mechanisms by which a user of the data marketplace 13008 may obtain access to one or more of the data elements.

[1465] The data payment allocation engine 13038 may allocate data marketplace payments. The data payment allocation engine 13038 may allocate data marketplace payments according to the value of a data stream 13010, the value of a contribution to a data stream 13010 and the like. This type of payment allocation may allow the data marketplace 13008 to allocate payments to data contributors, based on the value of the data contributions.

[1466] For example, contributors of data to a higher-value data stream 13010 may receive higher payments than contributors of data to lower-value data streams 13010. Similarly, data marketplace participants, for example IoT device manufacturers and system integrators, may be rated or ranked by the value of the data or the power of the configurations they provide and support.

[1467] The data marketplace 13008 may be a self-organizing data marketplace. A self-organizing data marketplace may self-organize using self-organization capabilities 13012. Self-organization capabilities 13012 may be learned, developed and optimized using artificial intelligence (AI) capabilities. AI capabilities may be provided by the machine learning 13014 capability, for example. Self-organization may occur via an expert system and may be based on the application of a model, one or more rules, or the like. Self-organization may occur via a neural network or deep learning system, such as by optimizing variations of the organization of the data pool over time based on feedback to one or more measures of success. Self-organization may occur by a

hybrid or combination of a rule-based system, model-based system, and neural network or other AI system. Various capabilities may be self-organized, such as how data elements are presented in the user interface of the marketplace, what data elements are presented, what data streams are obtained as inputs to the marketplace, how data elements are described, what metadata is provided with data elements, how data elements are stored (such as in a cache or other "hot" storage or in slower, but less expensive storage locations), where data elements are stored (such as in edge elements of a network), how data elements are combined, fused or multiplexed, or the like. Feedback to self-organization may include various metrics and measures of success, such as profit measures, yield measures, ratings (such as by users, purchasers, licensees, reviewers, and the like), indicators of interest (such as clickstream activity, time spent on a page, time spent reviewing elements and links to data elements), and others as described throughout this disclosure.

[1468] Data marketplace inputs 13056, data streams 13010 and data pools 13070 may be organized, based on metrics and measures of success of the data marketplace 13056. Data marketplace inputs 13056, data streams 13010 and data pools 13070 may be organized by the self-organization capability 13012, allowing the marketplace inputs 13056, data streams 13010, and data pools 13070 to be organized automatically, without requiring interaction by a user of the data marketplace 13008.

[1469] The metric and measure of success may also be used to configure the data brokering engine 13042 to execute a transaction among at least two marketplace participants. The machine learning engine 13014 may use the metric of success to configure the data brokering engine 13042 automatically, without requiring user intervention. The metric of success may also be used by a pricing engine, for example the marketplace value rating engine 13040, to set the price of one or more data elements within the data marketplace 13008.

[1470] In an exemplary and non-limiting embodiment, the self-organizing data marketplace may self-organize to determine which type of data streams 13010 are the most valuable and offer the most valuable and other data streams 13010 for sale. The calculation of data stream value may be performed by the marketplace value rating engine 13040.

[1471] In example embodiments, a policy automation system for a data collection system in an industrial environment may comprise: a policy input interface structured to receive policy inputs relating to definition of at least one parameter of at least one of a rule, a policy and a protocol, wherein the at least one parameter defines at least one of a configuration for a data collection device, an access policy for accessing data from the data collection device, and collection policy for collection of data by the device; and a policy automation engine for taking the inputs and automatically configuring and deploying at least one of the rule, the policy and the protocol within the system for data collection. In example embodiments, the at least one parameter may define at least one of an energy utilization policy, a cost-based policy, a data writing policy, and a data storage policy. The parameter may relate to a policy selected from among compliance, fault, configuration, accounting, provisioning and security policies for defining how devices are created, deployed and managed. The compliance policies may include data ownership policies. The data ownership policies may specify who owns data. The data ownership policies may specify how

owners may use data. The compliance policies may include data analysis policies. The data analysis policies may specify what data holders may access, how data holders may use data, and how data may be combined with other data by data holders. The compliance policies may include data use policies, data format policies, and the like. The data format policies may include standard data format policies, mandated data format policies. The compliance policies may include data transmission policies. The data transmission policies may include inter-jurisdictional transmission data transmission policies. The compliance policies may include data security policies, data privacy policies, information sharing policies, and the like. The data security policies may include at rest data security policies, transmitted data security policies, and the like. The information sharing policies may include policies specifying when information may be sold, when information may be shared, and the like. The compliance policies may include jurisdictional policies. The jurisdictional policies may include policies specifying who controls data. The jurisdictional policies may include policies specifying when data may be controlled. The jurisdictional policies may include policies specifying how data transmitted across boundaries is controlled.

[1472] In example embodiments, a policy automation system for a data collection system in an industrial environment may comprise: a policy automation engine for enabling configuration of a plurality of policies applicable to collection and utilization of data handled by a plurality of network connected devices deployed in a plurality of industrial environments, wherein the policy automation engine is hosted on information technology infrastructure elements that are located separately from the industrial environment, wherein upon configuration of a policy in the policy automation engine, the policy is automatically deployed across a plurality of devices in the plurality of industrial environments, wherein the policy sets configuration parameters relating to what data is collected by the data collection system and relating to access permissions for the collected data. The policies may include a plurality of policies selected among compliance, fault, configuration, accounting, provisioning and security policies for defining how devices are created, deployed and managed, and the plurality of policies communicatively coupled to policies. A policy input interface may be structured to receive policy inputs used as an input to at least one of a rule, policy and protocol definition, such as where the policy automation system a centralized source of policies for creating, deploying and managing policies for devices within an industrial environment.

[1473] In example embodiments, a policy automation system for a data collection system in an industrial environment may comprise: a policy automation engine for enabling configuration of a plurality of policies applicable to collection and utilization of data handled by a plurality of network connected devices deployed in a plurality of industrial environments, wherein the policy automation engine is hosted on information technology infrastructure elements that are located separately from the industrial environment, wherein upon configuration of a policy in the policy automation engine, the policy is automatically deployed across a plurality of devices in the plurality of industrial environments, wherein the policy sets configuration parameters relating to what data is collected by the data collection system and relating to access permissions for the collected data, wherein the policy automation system is communicatively coupled to a plurality of devices through a cloud network connection. The cloud network connection may be a privately-owned cloud

connection, a publicly provided cloud connection, a publicly provided cloud connection, the primary connection between the policy automation system and device, the primary connection between the policy automation system and device, an intranet cloud connection, connecting devices within a single enterprise, an extranet cloud connection, connecting devices among multiple enterprises, a secure cloud network connection, secured by a virtual private network (VPN) connection, and the like.

[1474] In example embodiments, a data marketplace for a data collection system in an industrial environment may comprise: an input interface structured to receive marketplace inputs; at least one of a data pool and a data stream to provide collected data within the marketplace; and data streams that include data from data pools. In example embodiments, at least one parameter of the marketplace may be automatically configured by a machine learning facility based on a metric of success of the marketplace. The inputs may include a plurality of data streams from a plurality of industrial data collectors. The data collectors may be multiplexing data collectors. The inputs may include consortia inputs. A consortium may be an existing consortium, a new consortium, a new consortium related to a data stream through a common interest, and the like. The metrics and measures of success may include profit measures, yield measures, ratings, indicators of interest, and the like. The ratings may include user ratings, purchaser ratings, licensee ratings, reviewer ratings, and the like. The indicators of interest may include clickstream activity, time spent on a page, time spent reviewing elements, links to data elements, and the like.

[1475] In example embodiments, a data marketplace for a data collection system in an industrial environment may comprise: an input system structured to receive a plurality of data inputs relating to data sensed from or about one or more industrial machines; at least one of a data pool and a data stream to provide collected data within the marketplace; and a self-organization system for organizing at least one of the data inputs and the data pools based on a metric of success of the marketplace. In example embodiments, the self-organization system may optimize variations of the organization of the data pool over time. The optimized variations may be based on feedback to one or more measures of success. The self-organization system may organize how data elements are presented in the user interface of the marketplace. The self-organization system may select what data elements are presented, what data streams are obtained as inputs to the marketplace, how data elements are described, what metadata is provided with data elements, a storage method for data elements, a location within a communication network for the storage elements (such as in edge elements of a network), a data element combination method, and the like. A storage method may include a cache or other "hot" storage method. A storage method may include slower, but less expensive storage locations. The data element combination method may be a data fusion method, a data multiplexing method, and the like. The self-organization system may receive feedback data, such as where feedback data includes success metrics and measures. Success metrics and measures may include profit measures, include yield measures, ratings, indicators of interest, and the like. Ratings include ratings that may be provided by users, purchasers, by licensees, reviewers. Success metrics and measures may include indicators of interest. Indicators of interest may include clickstream activity, time spent on a page activity, time spent reviewing elements, time spent

reviewing elements, links to data elements, and the like. The self-organization system may determine the value of data streams. The value of data streams may determine which data streams are offered for sale by the data marketplace. The ratings may include user ratings. The ratings may include purchaser ratings, licensee ratings, reviewer ratings, and the like.

[1476] In example embodiments, a data marketplace for a data collection system in an industrial environment may comprise: an input interface structured to receive data inputs from or about one or more of a plurality of industrial machines; at least one of a data pool and a data stream to provide collected data within the marketplace; and a rights management engine for managing permissions to access the data in the marketplace. In example embodiments, at least one parameter of the rights management engine may be automatically configured by a machine learning facility based on a metric of success of the marketplace. The rights management engine may assign rights to participants of the data marketplace. The rights may include business strategy and solution rights, liaison rights, marketing rights, security rights, technology rights, testbed rights, and the like. The metrics and measures of success may include profit measures, yield measures, ratings, and the like. The ratings may include user ratings, purchaser ratings, licensee ratings, reviewer ratings, and the like. The metrics and measures success may include indicators of interest, such as where interest includes clickstream activity, time spent on a page, time spent reviewing elements, and links to data elements.

[1477] In example embodiments, a data marketplace for a data collection system in an industrial environment may comprise: an input interface structured to receive data inputs from or about one or more of a plurality of industrial machines; at least one of a data pool and a data stream to provide collected data within the marketplace; and a data brokering engine configured to execute a data transaction among at least two marketplace participants. In example embodiments, at least one parameter of the data brokering engine may be automatically configured by a machine learning facility based on a metric of success of the marketplace. A data transaction input may include a marketplace value rating. A marketplace value rating may be assigned to a marketplace participant. A marketplace value rating may be assigned to a marketplace participant is assigned based on the value of input provided by the participant to the marketplace. A data transaction may be a trade transaction, a sale transaction, is a payment transaction, and the like. The metrics and measures of success may include profit measures, yield measures, ratings, and the like. The ratings may include user ratings. The ratings may include purchaser ratings, licensee ratings, reviewer ratings, and the like. The metrics and measures success may include indicators of interest. The indicators of interest may include clickstream activity, time spent on a page, include time spent reviewing elements, links to data elements, and the like.

[1478] In example embodiments, a data marketplace for a data collection system in an industrial environment may comprise: an input interface structured to receive data inputs from or about one or more of a plurality of industrial machines; at least one of a data pool and a data stream to provide collected data within the marketplace; and a pricing engine for setting a price for at least one data element within the marketplace. In example embodiments, pricing may be automatically configured for the pricing engine by a machine learning facility based on a metric of success of the

marketplace. The metrics and measures of success may include profit measures, yield measures, ratings, and the like. The ratings may include user ratings. The ratings may include purchaser ratings, licensee ratings, reviewer ratings, and the like. The metrics and measures success may include indicators of interest. The indicators of interest may include clickstream activity, time spent on a page, include time spent reviewing elements, links to data elements, and the like.

[1479] In example embodiments, a data marketplace for a data collection system in an industrial environment may comprise: an input interface structured to receive data inputs from or about one or more of a plurality of industrial machines; at least one of a data pool and a data stream to provide collected data within the marketplace; and a user interface for presenting a data element and at least one mechanism by which a party using the marketplace can obtain access to the at least one data stream or data pool. In example embodiments, pricing may be automatically configured for the pricing engine by a machine learning facility based on a metric of success of the marketplace. The metrics and measures of success may include profit measures, yield measures, ratings, and the like. The ratings may include user ratings. The ratings may include purchaser ratings, licensee ratings, reviewer ratings, and the like. The metrics and measures success may include indicators of interest. The indicators of interest may include clickstream activity, time spent on a page, include time spent reviewing elements, links to data elements, and the like.

[1480] In example embodiments, a data collection system in an industrial environment may comprise: a policy automation system for a data collection system in an industrial environment, comprising: a plurality of rules selected among roles, permissions and uses, the plurality of rules communicatively coupled to policies, protocols, and policy inputs; a plurality of policies selected among compliance, fault, configuration, accounting, provisioning, and security policies for defining how devices are created, deployed and managed, the plurality of policies communicatively coupled to policies, protocols and policy inputs and a policy input interface structured to receive policy inputs used as an input to at least one of a rule, policy and protocol definition.

[1481] In example embodiments, a data marketplace may comprise: an input interface structured to receive marketplace inputs; a plurality of data pools to store collected data, including marketplace inputs and make collected data available for use by the marketplace; and data streams that include data from data pools.

[1482] Intelligent industrial equipment and systems may be configured in various networks, including self-forming networks, private networks, Internet-based networks, and the like. One or more of the smart heating systems may incorporate hydrogen production, storage, and use may be configured as nodes in such a network. In example embodiments, a smart heating system may be configured with one or more network ports, such as a wireless network port that facilitates connection through Wi-Fi and other wired and/or wireless communication protocols as described. The smart heating system includes a smart hydrogen production system and a smart hydrogen storage system, which may be configured individually or as an integral system connected as one or more nodes in a network of industrial equipment and systems. By way of this example, a smart heating system may be disposed in an on-site industrial equipment operations center, such as a

portable trailer equipped with communication capabilities and the like. Such deployed smart heating system may be configured, manually, automatically, or semi-automatically to join a network of devices, such as industrial data collection, control, and monitoring nodes and participate in network management, communication, data collection, data monitoring, control, and the like.

[1483] In another example of a smart heating system participating in a network of industrial equipment monitoring, control, and data collection devices in that a plurality of the smart heating systems may be configured into a smart heating system sub-network. In example embodiments, data generated by the sub-network of devices may be communicated over the network of industrial equipment using the methods and systems described herein.

[1484] In example embodiments, the smart heating system may participate in a network of industrial equipment as described herein. By way of this example, one or more of the smart heating systems, as depicted in Figure 136, may be configured as an IoT device, such as IoT device 13500 and the like described herein. In example embodiments, the smart heating system 13502 may communicate through an access point, over a mobile ad hoc network or mechanism for connectivity described herein for devices and systems elements and/or through network elements described herein.

[1485] In example embodiments, one or more smart heating systems may incorporate, integrate, use, or connect with facilities, platforms, modules, and the like that may enable the smart heating system to perform functions such as analytics, self-organizing storage, data collection and the like that may improve data collection, deploy increased intelligence, and the like. Various data analysis techniques, such as machine pattern recognition of data, collection, generation, storage, and communication of fusion data from analog industrial sensors, multi-sensor data collection and multiplexing, self-organizing data pools, self-organizing swarm of industrial data collectors, and others described herein may be embodied in, enabled by, used in combination with, and derived from data collected by one or more of the smart heating systems.

[1486] In example embodiments, a smart heating system may be configured with local data collection capabilities for obtaining long blocks of data (i.e., long duration of data acquisition), such as from a plurality of sensors, at a single relatively high-sampling rate as opposed to multiple sets of data taken at different sampling rates. By way of this example, the local data collection capabilities may include planning data acquisition routes based on historical templates and the like. In example embodiments, the local data collection capabilities may include managing data collection bands, such as bands that define a specific frequency band and at least one of a group of spectral peaks, true-peak level, crest factor and the like.

[1487] In example embodiments, one or more smart heating systems may participate as a self-organizing swarm of IoT devices that may facilitate industrial data collection. The smart heating systems may organize with other smart heating systems, IoT devices, industrial data collectors, and the like to organize among themselves to optimize data collection based on the capabilities and conditions of the smart heating system and needs to sense, record, and acquire information from and around the smart heating systems. In example embodiments, one or more smart heating systems may be configured with processing intelligence and capabilities that may facilitate

coordinating with other members, devices, or the like of the swarm. In example embodiments, a smart heating system member of the swarm may track information about what other smart heating systems in a swarm are handling and collecting to facilitate allocating data collection activities, data storage, data processing and data publishing among the swarm members.

[1488] In example embodiments, a plurality of smart heating systems may be configured with distinct burners but may share a common hydrogen production system and/or a common hydrogen storage system. In example embodiments, the plurality of smart heating systems may coordinate data collection associated with the common hydrogen production and/or storage systems so that data collection is not unnecessarily duplicated by multiple smart heating systems. In example embodiments, a smart heating system that may be consuming hydrogen may perform the hydrogen production and/or storage data collection so that as smart heating system may prepare to consume hydrogen, they coordinate with other smart heating systems to ensure that their consumption is tracked, even if another smart heating system performs the data collection, handling, and the like. In example embodiments, smart heating systems in a swarm may communicate among each other to determine which smart heating system will perform hydrogen consumption data collection and processing when each smart heating system prepares to stop consumption of hydrogen, such as when heating, cooking, or other use of the heat is nearing completion and the like. By way of this example when a plurality of smart heating systems is actively consuming hydrogen, data collection may be performed by a first smart heating system, data analytics may be performed by a second smart heating system, and data and data analytics recording or reporting may be performed by a third smart heating system. By allocating certain data collection, processing, storage, and reporting functions to different smart heating systems, certain smart heating systems with sufficient storage, processing bandwidth, communication bandwidth, available energy supply and the like may be allocated an appropriate role. When a smart heating system is nearing an end of its heating time, cooking time, or the like, it may signal to the swarm that it will be going into power conservation mode soon and, therefore, it may not be allocated to perform data analysis or the like that would need to be interrupted by the power conservation mode.

[1489] In example embodiments, another benefit of using a swarm of smart heating systems as disclosed herein is that data storage capabilities of the swarm may be utilized to store more information than could be stored on a single smart heating system by sharing the role of storing data for the swarm.

[1490] In example embodiments, the self-organizing swarm of smart heating systems includes one of the systems being designated as a master swarm participant that may facilitate decision making regarding the allocation of resources of the individual smart heating systems in the swarm for data collection, processing, storage, reporting and the like activities.

[1491] In example embodiments, the methods and systems of self-organizing swarm of industrial data collectors may include a plurality of additional functions, capabilities, features, operating modes, and the like described herein. In example embodiments, a smart heating system may be configured to perform any or all of these additional features, capabilities, functions, and the like without limitation.

[1492] The server may provide an interface to other devices including, without limitation, clients, other servers, printers, database servers, print servers, file servers, communication servers, distributed servers, social networks, and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking of some or all of these devices may facilitate parallel processing of a program or method at one or more locations without deviating from the scope of the disclosure. In addition, any of the devices attached to the server through an interface may include at least one storage medium capable of storing methods, programs, code, and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

[1493] The software program may be associated with a client that may include a file client, print client, domain client, internet client, intranet client, and other variants such as secondary client, host client, distributed client, and the like. The client may include one or more of memories, processors, computer readable transitory and/or non-transitory media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other clients, servers, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the client. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the client.

[1494] The client may provide an interface to other devices including, without limitation, servers, other clients, printers, database servers, print servers, file servers, communication servers, distributed servers, and the like. Additionally, this coupling and/or connection may facilitate remote execution of a program across the network. The networking of some or all of these devices may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the client through an interface may include at least one storage medium capable of storing methods, programs, applications, code, and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

[1495] FIG. 286 provides a schematic view of an architecture, its components and functional relationships for an industrial Internet of Things solution. A data handling platform 13700 may include a set of data handling layers, such as those described in various embodiments throughout this disclosure and the documents incorporated by referenced herein, such as intelligent storage systems and capabilities 13724, monitoring and collection systems and capabilities 13728, and processing and intelligence capabilities 13730, which may serve a set of applications 13732, such as where various capabilities, microservices, and the like of the platform 13700 may service multiple applications 13732 in a unified or integrated way. The platform 13700 may be deployed in a cloud computing environment, such as on cloud computing infrastructure and services, and may, such as through the monitoring and collection systems and capabilities 13728, connect to an industrial environment 13704, where an edge system 13718 may provide connectivity (such as

using any of the network and/or software systems, services, protocols, or capabilities described throughout this disclosure and the documents incorporated by reference herein, or as would be understood by one skilled in the art, such as cellular connectivity (including 5G capabilities), Wi-Fi, BLUETOOTH™, and other network protocols, and application programming interfaces, ports, connectors, brokers, and other software systems, among many others), computation (such as processing of data, signal processing, data transformation, data normalization, or the like) and intelligence (such as applying decision rules or models, computing and operating on inputs to produce analytics, alerts, reports and/or control instructions, applying one or more artificial intelligence systems (such as machine learning systems, neural networks, expert systems, deep learning systems, or other systems disclosed throughout this disclosure or in the document incorporated by reference here).

[1496] In example embodiments, the platform 13700 may generate, host, integrate, link to, include, integrate, or otherwise interact with a set of industrial entity digital twins 13734, which may comprise digital representations or replicas of real world states of a set of industrial entities 13736, such as workers 13712, fixed assets 13712 (such as machines, systems, devices, fixtures, and the like), infrastructure 13710 (such as floors, walls, ceilings, loading docks, foundations, and many others), and moving assets 13708 (such as vehicles, forklifts, autonomous vehicles, drones, assembly lines, fans, rotors, turbines, pumps, valves, fluids, and many others), among many others that reside in or about an industrial environment 13704.

[1497] In example embodiments, the edge system 13718 may interface with each of a mobile data collector 13702 (such as having any of the capabilities described throughout this disclosure and the documents incorporated by reference herein, such as having onboard intelligence (such as for optimizing data storage and processing, power utilization, or selection of a set of sensor inputs among various available inputs, such as having a cross-point switch or similar facility for selecting and routing a subset of sensor channels, such as having an RFID reader or other reader for taking asset tag and similar data from entities 13736, or such as having capabilities to connect to and read from sensors 13722 and/or onboard diagnostic systems, buses, and other systems that are integrated with or into the entities 13736, among many others) and a simultaneous location and mapping (SLAM) system 13714 (such as for precisely determining locations of entities 13736 within a space and mapping those entities to the locations, such as by representing the entities 13736 in a point cloud that represents the results of scanning the environment 13704 or part thereof with LIDAR, ultrasound, sonar, X-ray, magnetic resonance imaging, infrared, deep infrared, or other scanning technology that is capable of providing a representation of the entities 13736 within the environment 13704. In one illustrative embodiment, a scan represents the entities 13736 as a point cloud of data points collected by a LIDAR-based SLAM system 13714. In example embodiments, the mobile data collector 13702 and the SLAM system 13714 are integrated or linked, so that locations, positions, orientations, or the like, of the points of collection of data by the mobile data collector 13702 are automatically registered with or by the SLAM system 114, such that a unified data set is provided to the edge system 13718 for further communication, computation, or processing. For example, a set of vibration data readings made by a 3-axis mobile data collector

13702 may be registered to particular locations of a mapped point cloud of data created by or in the SLAM system 13714, so that vibration information can be linked to those parts of the point cloud and subsequently linked to a machine or other entity 13730 represented by that part of the point cloud or other mapping systems.

[1498] In example embodiments, the SLAM system 13714 and the mobile data collector 13702 are integrated into a single portable device, allowing a data collection route to be performed (such as by a worker, drone, or autonomous vehicle) as the space is mapped by the SLAM system 13714. This may thus comprise a simultaneous location, mapping and data collection system (SLAMDC) 13740. In example embodiments, the mobile data collector 13702 may collect from sensors 13722 that are included in or integrated with the data collector 13702 (such as onboard triaxial vibration sensors, ultrasound sensors, acoustic sensors, heat sensors, or many others, including any of the types of sensors disclosed herein or in the documents incorporated herein by reference). In example embodiments, the mobile data collector 13702 may collect from sensors 13722 that are disposed in or around the environment 13704, such as cameras, analog sensors, digital sensors, or many others. In example embodiments, a data collector, such as a worker, drone, or autonomous vehicle, may be instructed, such as by onboard intelligence, intelligence of the edge system 13718, or an intelligent system 13730 of the platform 13700, to place additional sensors in, on, or in proximity to a set of industrial entities 13736, such as where intelligence indicates a benefit to additional gathering of information, such as where a problem is detected or predicted.

[1499] In example embodiments, the edge system 13718 may include, link or connect to, integrate with, or be integrated into a control system 13742, such as for providing control for one or more industrial entities 13736, such as controlling a machine in a factory (such as a CNC machine, additive manufacturing machine, energy system (e.g., a generator or turbine), an assembly line, or the like), controlling a workflow (such as a production workflow, an inspection workflow, a data collection workflow, a maintenance workflow, a servicing workflow, or the like), or controlling sub-systems, systems, or operations of an entire factory or set of factories. Processing, computation and intelligence capabilities of the edge system 13718 (and by virtue of connectivity between the edge system 13718 and the platform 13700, processing, computation and intelligence capabilities 13730 of the platform 13700) may thus benefit from input from a set of control systems 13742 and may provide inputs to (including control signals for) the set of control systems 13742. Data from the mobile data collector 13702 (including from sensors 13722, onboard information of the entities 13736, and other information), from the edge system 13718, from the SLAM system 13714, from a combined SLAMDC system 13740, from one or more applications 13732, or from the platform 13700 (including any of the layers there), may be represented in a set of industrial digital twins 13734. For example, an industrial digital twin 13734 may show a point cloud view of a mapped industrial environment (which, in example embodiments, may be augmented, such as using 3D mapping, AR or VR systems) with relevant data collection elements presented in the point cloud view along with the point cloud. Many examples are available, such as highlighting (such as by color or motion) in the digital twin 13734, areas of the point cloud where systems are vibrating in a way that is out of the normal range (such as where severity units, as discussed elsewhere herein,

exceed a threshold). Industrial entity digital twins 13734 may include, link or connect to, or integrate with a variety of interfaces and dashboards 13738, such as ones configured for specific workflows, roles, and users. For example, dashboards and interfaces may be configured for workers who will interact with specific machines (such as where the digital twin is used for training, workflow guidance, diagnosis of problems, and the like); for managers of operations on a factory floor (such as where a digital twin 13734 displays a layout of machines on the floor, patterns of traffic (e.g., moving assets. 13708 and workers 13712) involved in workflows, status information for workers, machines, processes, or the like (including operational status, maintenance status, inspection status, and the like), analytic information (such as indicating metrics about operations, about potential problems, or the like); for inspectors (such as where the digital twin 13734 represents areas that are indicated by data collectors 13702 to require or benefit from additional inspection (e.g., where the inspector can check off items that have already been inspected or highlight items for further inspection by interacting with them in a digital twin interface or dashboard 13738); for maintenance and service workers (such as where a digital twin 13734 highlights locations of items requiring maintenance in a schematic view and guides the service workers to the right location and/or machine, then presents (such as in a different view) information and guidance on how to undertake the service or maintenance, ranging from a checklist or workflow to a virtual, mixed or augmented reality training or guidance session that can be presented at the machine); for front office managers (such as finance professionals who can be presented financial information, such as ROI metrics, output metrics, cost metrics, and the like (including current status and predictions), legal personnel (such as where a digital twin 13734 may present compliance information, highlight legal risks (such as safety violations or instances where status information about operations indicates a likelihood that the company may breach a contract (such as by failing to produce an output that is required by a contract) and the like), inventory managers, procurement personnel, and the like; and for executives, such as CEOs, CTOs, COOs, CIOs, CDOs, CMOs, and the like, who may interact with digital twins 13734 that represent whole factories, or sets of factories, such as to identify risks and opportunities that may involve understanding interactions of elements and/or contributions of elements involving industrial entities 13736 to overall operations of an enterprise, to its strategies, or the like.

[1500] In various embodiments, the interfaces and dashboards 13738 may display sensor information collected from sensors 13722, from mobile data collectors 13702, from SLAM systems 13714 (or combined SLAMDC systems 13740); mapping information from a SLAM system 13714 or SLAMDC system 13740; representations of shapes and placements of entities 13736 (such as point clouds, CAD drawings, photographs, 3D representations, blueprints, or abstract representations (such as topologies or hierarchies showing relationships); representations of calculations, metrics, computations, statistics, analytics and the like (such as computed by the edge system 13718, processing and intelligence system 13730, or other system); state or status information (such as indicating operational states or status of workflows involving industrial entities 13736), or the like.

[1501] Information elements from the industrial environment 13704 or about industrial entities 13736 can be presented in overlays (e.g., where metrics or symbols are presented on top of a point cloud, a photo, or a 3D representation of a unit in a 3D interface), in native form (such as where a point cloud is represented), in 3D visualizations (such as where the interface handles elements as 3D geometric elements), and the like.

[1502] Interfaces and dashboards 13738 may include graphical interfaces (such as for laptops, tablets and mobile devices), touch screen interfaces, voice-activated interfaces, augmented reality interfaces, virtual reality interfaces, mixed reality interfaces, application programming interfaces (APIs), and the like.

[1503] Digital twins 13734 may be of various types, such as component digital twins represent an individual part of component; machine digital twins that represent an entire machine; system digital twins that represent a system involving multiple components, parts, machines or the like and their interactions; worker digital twins that represent one or more attributes or states of a set of workers 13712; arrangement digital twins that represent the layout or arrangement of entities 13736 (such as, without limitation, the arrangement of components, assets, machines, workers, or other elements on a factory floor); augmented, virtual and/or mixed reality digital twins that provide a realistic experience for a user, such as simulating or mimicking interaction with an asset, another worker, a workflow, or the like (such as for training a worker or group of workers how to operate or undertake maintenance on a machine or system, how to undertake a workflow involving a machine or system, or the like); abstract digital twins (such as ones that represent elements and relationships, such as in topologies, hierarchies, flow diagrams, or the like), and others.

[1504] In example embodiments, interfaces and dashboards 13738 may be provided that facilitate drilling down and/or zooming up in a digital twin 13734 (whether under user control or by automation, such as based on an understanding of status information, contextual information, user interactions, or other factors), such as to obtain a more detailed view of a component of a larger view (e.g., to see a specific part of a machine in an exploded view); to move up to a wider view that encompasses more components and/or their interactions; to obtain additional information (such as to see additional metrics related to a metric represented in a digital twin 13734, more granular data, source data that was used to determine a metric, or the like); and the like.

[1505] In example embodiments, interfaces and dashboard 13738 may be configured to facilitate switching between views or types of digital twin of the same entities 13736 (whether under user control or by automation, such as based on an understanding of status information, contextual information, user interactions, or other factors involving the digital twin 13734). For example, a user may switch from an overall schematic view that represents current status information for the machines and workflows on a factory floor to a 3D view that shows a realistic representation of one of the machines (such as one that has been highlighted as having an issue, such as where a data collector 13702 has determined that it is operating outside normal parameters for temperature, vibration, pressure, or the like).

[1506] In various embodiments an end-to-end system is provided, where an industrial digital twin 13734 maintains an ongoing or periodically updated data connection, via one or more layers of the

platform 100, through connectivity to an edge system, to a mobile data collector 13702, SLAM system 13714 and/or SLAMDC system 13740, such that the industrial digital twin 13734 provides real-time, or periodically updated, information about the current attributes, states, status, or the like of the entities 13736 in an industrial environment 13704. This may include, as noted above, representing sensor data from sensors 13722, onboard data from entities 13736, control information from control systems 13742, various data collected by data collectors 13702, mapping information, information computed by edge intelligence of an edge system 13718 and/or processing and intelligence system 13730, and the like, such that a manager, executive, or other users can have highly interactive visualization of and interaction with the elements under the user's authority, or otherwise of interest to the user.

[1507] In example embodiments, analytics derived from data collection by a mobile data collector 13702 and/or from sensors 13722, control systems 13742 and/or onboard sensing or diagnostics of industrial entities 13736 may be computed by the edge system 13718 and/or the processing and intelligence system 13730 may include a metric that indicates, based on current information from these various data sources, and optionally based on historical data from outcomes involving similar entities 13736, the probability of an unscheduled shutdown during a period of time. The unscheduled shutdown metric may be calculated for various entities 13736, such as for a machine, a system, a workflow, a factory, or a set of factories, and it may be represented in an industrial digital twin 13734, such as by representing the metric as an overlay element on a digital twin that provides a schematic of a factory floor.

[1508] In example embodiments, contributing component factors to the probability of an unscheduled shutdown of an industrial entity 13736, a workflow, or an operation and the like may be analyzed and represented in an interface or dashboard 13738 of an industrial digital twin 13734. These component factors may include the probability of occurrence of known failure modes of components or machines (such as calculated by predictive maintenance models, such as ones that use physical models, historical models, historical models, or the like), such as failures based on mechanical stress, overloading, wear and tear, problems with bearings, problems with couplings, out-of-balance states of rotating components, overheating, freezing, excess viscosity, lubrication problems, clogging, cavitation, vacuum failures, leaks, low fluid levels, low pressure levels, electrical failures, power failures, failures of component supply, absence of tools, absence of component parts, broken parts, shutdowns of other entities 13736, traffic congestion, information technology problems, computation errors, cyberattacks, and many others.

[1509] In example embodiments, unscheduled shutdown probability may be determined by a prediction machine, such as a neural network, such as one that is trained on a historical data set of failures. In example embodiments, unscheduled shutdown probability for an entity may be determined by a combination of a model-based approach and a neural network, such as where a neural network determines a probability of a specific type of failure and/or of a specific part of a system, and that probability is used in a model to compute a probability of shutdown of a system in which that type of failure or specific part is involved, or vice versa.

[1510] In example embodiments, unscheduled shutdown probabilities may be computed at the edge system 13718, by the processing and intelligence capabilities 13730 of the platform 100, or by a combination of those or other intelligence systems. Unscheduled shutdown probability metrics may be represented in a set of industrial digital twins 13734, such as providing managers, maintenance workers, executives, inspectors, and others a visual indication of the overall risk of an unscheduled shutdown, as well as visual indicators of the component elements or entities 13736 that are at risk, or that are contributing to increases in the probability of an unscheduled shutdown of a factory, plant, system, process, line, machine, workflow, or the like. This may allow managers and executives to drill down, obtain further information, and undertake actions that reduce the risk. As one illustrative example, an executive may be presented with a view of a set of factories, with one factory being represented in an industrial digital twin 13734 in a different color (such as bright red) based on that factory having a probability of unscheduled shutdown that exceeds a threshold (or simply that it has the highest probability among a set of factories). This may direct the attention of the executive to that factory, thereby leading to further insight into operational choices that would have been missed if the executive were merely presented with raw data, a spreadsheet, or the like where the unscheduled shutdown probability would need to be calculated, inferred, or the like. Similarly, a factory manager for the highlighted factory may have an industrial digital twin 13734 that presents the probabilities of unscheduled shutdown of various component machines and processes; for example, a pump that is maintaining a vacuum of a critical semiconductor production process for the factory (or a biologics production process, or the like) may be identified as having a high risk of failure, such as based on vibration analysis that indicates cavitation, in combination with other data sources, such as ones indicating the age of the pump and its maintenance and operating history. The pump may be highlighted in the industrial digital twin 13734, such as in a view configured for the factory manager, such as by highlighted the pump in a bright color and by animating the pump with movement (such as shaking a visual element) that indicates a vibration problem is the likely contributor to the risk of unscheduled shutdown of the pump (which cascades to a failure of the vacuum, the failure of the critical production process, and the shutdown of the entire factory). As a result of attention being directed by the digital twin by visual cues (as compared to a spreadsheet or raw data output), the factory manager may direct (including by interacting with the pump in the digital twin, such as by touching it) attention to the pump for maintenance or replacement. An instruction or message provided by one user (such as the factory manager or executive) may result in a message, or highlighting, in a different digital twin 13734 or user interface or dashboard 13738 that is configured for another user. For example, the pump, if flagged by the factory manager in a view of the factory, may appear in a service worker's digital twin 13734, such as showing a route to the pump and subsequently switching to a view that guides the worker through inspection, maintenance, service, and/or replacement. Thus, a set of digital twins 13734 may highlight unscheduled shutdown risks based on real-time or periodic connection through edge intelligence to data collection systems and facilitate workflows (enabled within the digital twins) by which attention is directed for various workers (by highlighting visual elements)

to issues that they can address, optionally with guidance and instruction from additional views of the set of digital twins.

[1511] In example embodiments, the end-to-end real time or periodic connection between a set of industrial digital twins 13734 through the platform 13700, the edge system 13718, control systems 13742, data collectors 13702, SLAM systems 13714, SLAMDC systems 13740 and sensors 13722 to industrial entities 13736 and their various onboard sensors, data collection systems, diagnostic systems, buses, and the like may facilitate control over the various elements of these systems via manipulation of elements in interfaces and dashboards 13738 of the digital twins 13734, including ones that are linked to, included in, or integrated with one or more applications 13732, such as via APIs. For example, manipulating an element of an industrial digital twin 13734 may be used to configure or modify data collection by a mobile data collector 13702, such as by causing the mobile data collector 13702 to switch channels (such as where multiple sensor channels are available, and (such as via a cross-point switch) the data collector 13702 is instructed to switch from, for example, collecting a single axis vibration channel, temperature and pressure to collecting three-axis vibration data. This may occur for example, if a manager sees a potential vibration problem in a digital twin 13734 of a machine and touches the element for a drill down, which may automatically, or under user control, switch the data collection mode to provide different sensor data, more granular data (such as by collecting data at much shorter time intervals or in a streaming format, or the like). As another example, manipulating a user interface element or dashboard element 13738 or providing an instruction via an API to a digital twin 13734 may configure or modify configuration of intelligence or computation capabilities, such as of an edge system 13718, a processing and intelligence system 13730 of the platform 13700, or other intelligence system; for example, a user (or the system, under automated control), may reconfigure the edge system to access different data sources, such as by pruning data sources that appear to have little influence or adding new data sources that may improve outcomes, such as ones involving classification activities, prediction activities, and or control activities. For example, a predictive maintenance system (or multiple such systems) may exist for a factory. When the factory is scanned to produce a point cloud that represents various physical entities in the environment, such as during a data collection and mapping route of a SLAMDC system 13740, and the factory appears on the industrial digital twin 13734 of a user, the user may be presented with a set of additional data sources available for that factory, including the predictive maintenance data, and the user may select the data source and link it (such as by dragging and dropping it) to a part of the digital twin (e.g., where a point cloud represents a machine at a given location), resulting in the predictive maintenance data being fed as a data source to any intelligence systems that operate on that machine. Whether to facilitate augmenting intelligence systems as in this example, or for other purposes, the platform 13700 may facilitate connection of the end-to-end industrial digital twin system 13734 (and the elements that exchange information with it and/or are controlled by it) with other information technology systems of an enterprise, such as by linking to, providing inputs to, taking puts from, and/or integrating with those other systems, which may include, without limitation, enterprise resource planning systems, control systems, predictive maintenance systems,

inventory management systems, procurement systems, inspection systems, compliance systems, quality control systems, operations planning systems, and many others.

[1512] In example embodiments, manipulating a user interface element or dashboard element 13738 or providing an instruction via an API to a digital twin 13734 may configure or modify configuration of a control system 13742 or provide a control signal to a control system 13736, such that the digital twin provides a direct control interface to one or more industrial entities 13736.

[1513] In example embodiments, an industrial digital twin 13734 and related end-to-end system of data collection and intelligence may be used in connection with support of a service ecosystem, such as one where maintenance and service activities of the types disclosed throughout this disclosure and the documents incorporated by reference are supported, such as where an understanding of maintenance and service needs, in particular where intelligence indicates an elevated probability of unscheduled shutdown of an important entity 13736, is represented in a set of industrial digital twins 13734 configured for use by the users and applications (including ones that provide robotic process automation) involved in a service ecosystem, such as ones involved in identifying risks, flagging service issues, identifying and ordering necessary parts, tools, or components, identifying capable workers with necessary expertise, scheduling workers, parts, components and the like, scheduling necessary shutdowns of dependent processes and operations, routing workers and assets to service locations (outside and within the floor of a factory or plant), guiding workers (including automated workers) through procedures and protocols, prompting data collection and reporting, and many others. This support includes providing real-time and/or periodic updating from data collection, providing visualization of elements, with zooming, drilling down, switching views and the like (automatically and/or under user control), allowing interactions to obtain or configure intelligence and/or control, and other capabilities noted throughout this disclosure and in the documents incorporated herein by reference.

[1514] FIG. 287 is a schematic illustrating an industrial setting 28720 at which a sensor kit 28700 has been installed. In example embodiments, the sensor kit 28700 may refer to a fully deployable, purpose-configured industrial IoT system that is provided in a unified kit and is ready for deployment in the industrial setting 28720 by a consumer entity (e.g., owner or operator of an industrial setting 28720). In example embodiments, the sensor kit 28700 allows the owner or operator to install and deploy the sensor kit with no or minimal configuration (e.g., setting user permissions, setting passwords, and/or setting notification and/or display preferences). The term "sensor kit" 28700 may refer to a set of devices that are installed in an industrial setting 28720 (e.g., a factory, a mine, an oil field, an oil pipeline, a refinery, a commercial kitchen, an industrial complex, a storage facility, a building site, and the like). The collection of devices comprising the sensor kit 28700 includes a set of one or more internet of things (IoT) sensors 28702 and a set of one or more edge devices 28704. For purposes of discussion, references to "sensors" or "sensor devices" should be understood to mean IoT sensors, unless specifically stated otherwise.

[1515] In example embodiments, the sensor kit 28700 includes a set of IoT sensors 28702 that are configured for deployment in, on, or around an industrial component, a type of an industrial component (e.g., a turbine, a generator, a fan, a pump, a valve, an assembly line, a pipe or pipeline,

a food inspection line, a server rack, and the like), an industrial setting 28720, and/or a type of industrial setting 28720 (e.g., indoor, outdoor, manufacturing, mining, drilling, resource extraction, underground, underwater, and the like) and a set of edge devices capable of handling inputs from the sensors and providing network-based communications. In example embodiments, an edge device 28704 may include or may communicate with a local data processing system (e.g., a device configured to compress sensor data, filter sensor data, analyze sensor data, issue notifications based on sensor data and the like) capable of providing local outputs, such as of signals and of analytic results that result from local processing. In example embodiments, the edge device 28704 may include or may communicate with a communication system (e.g., a Wi-Fi chipset, a cellular chipset, a satellite transceiver, cognitive radio, one or more BLUETOOTH™ chips and/or other networking device) that is capable of communicating data (e.g., raw and/or processed sensor data, notifications, command instructions, etc.) within and outside the industrial environment. In example embodiments, the communication system is configured to operate without reliance on the main data or communication networks of an industrial setting 28720. In example embodiments, the communication system is provided with security capabilities and instructions that maintain complete physical and data separation from the main data or communication networks of an industrial setting 28720. For example, in example embodiments, BLUETOOTH™-enabled edge devices may be configured to permit pairing only with pre-registered components of a kit, rather than with other BLUETOOTH™-enabled devices in an industrial setting 28720.

[1516] In example embodiments, an IoT sensor 28702 is a sensor device that is configured to collect sensor data and to communicate sensor data to another device using at least one communication protocol. In example embodiments, IoT sensors 28702 are configured for deployment in, on, or around a defined type of an industrial entity. The term industrial entity may refer to any object that may be monitored in an industrial setting 28720. In example embodiments, industrial entities may include industrial components (e.g., a turbine, a generator, a fan, a pump, a valve, an assembly line, a pipe or pipeline, a food inspection line, a server rack, and the like). In example embodiments, industrial entities may include organisms that are associated with an industrial setting 28720 (e.g., humans working in the industrial setting 28720 or livestock being monitored in the industrial setting 28720). Depending on the intended use, setting, or purpose of the sensor kit 28700, the configuration and form factor of an IoT sensor 28702 will vary. Examples of different types of sensors include: vibration sensors, inertial sensors, temperature sensors, humidity sensors, motion sensors, LIDAR sensors, smoke/fire sensors, current sensors, pressure sensors, pH sensors, light sensors, radiation sensors, and the like.

[1517] In example embodiments, an edge device 28704 may be a computing device configured to receive sensor data from the one or more IoT sensors 28702 and perform one or more edge-related processes relating to the sensor data. An edge-related process may refer to a process that is performed at an edge device 28704 in order to store the sensor data, reduce bandwidth on a communication network, and/or reduce the computational resources required at a backend system. Examples of edge processes can include data filtering, signal filtering, data processing, compression, encoding, quick-predictions, quick-notifications, emergency alarming, and the like.

[1518] In example embodiments, a sensor kit 28700 is pre-configured such that the devices (e.g., sensors 28702, edge devices 28704, collection devices, gateways, etc.) within the sensor kit 28700 are configured to communicate with one another via a sensor kit network without a user having to configure the sensor kit network. A sensor kit network may refer to a closed communication network that is established between the various devices of the sensor kit and that utilizes two or more different communication protocols and/or communication mediums to enable communication of data between the devices and to a broader communication network, such as a public communication network 28790 (e.g., the Internet, a satellite network, and/or one or more cellular networks). For example, while some devices in a sensor kit network may communicate using a BLUETOOTH™ communication protocol, other devices may communicate with one another using a near-field communication protocol, a ZIGBEE® protocol, and/or a Wi-Fi communication protocol. In some implementations, a sensor kit 28700 may be configured to establish a mesh network having various devices acting as routing nodes within the sensor kit network. For example, sensors 28702 may be configured to collect data and transmit the collected data to the edge device 28704 via the sensor kit network, but may also be configured to receive and route data packets from other sensors 28702 within the sensor kit network towards an edge device 28704.

[1519] In example embodiments, a sensor kit network may include additional types of devices. In example embodiments, a sensor kit 28700 may include one or more collection devices (not shown in FIG. 138) that act as routing nodes in the sensor network, such that the collection devices may be part of a mesh network. In example embodiments, a sensor kit 28700 may include a gateway device (not shown in FIG. 138) that enable communication with a broader network, whereby the gateway device may communicate with the edge device 28704 over a wired or wireless communication medium in industrial settings 28720 that would prevent an edge device 28704 from communicating with the public communication network 28790 (e.g., in a factory having very thick concrete walls). Embodiments of the sensor kit 28700 may include additional devices without departing from the scope of the disclosure.

[1520] In example embodiments, the sensor kit 28700 is configured to communicate with a backend system 28750 via a communication network, such as the public communication network 28790. In example embodiments, the backend system 28750 is configured to receive sensor data from a sensor kit 28700 and to perform one or more backend operations on the received sensor data. Examples of backend operations may include storing the sensor data in a database, performing analytics tasks on the sensor data, providing the results of the analytics and/or visualizations of the sensor data to a user via a portal and/or a dashboard, training one or more machine-learned models using the sensor data, determining predictions and/or classifications relating to the operation of the industrial setting 28720 and/or industrial devices of the industrial setting 28720 based on the sensor data, controlling an aspect and/or an industrial device of the industrial setting 28720 based on the predictions and/or classifications, issuing notifications to the user via the portal and/or the dashboard based on the predictions and/or classifications, and the like.

[1521] It is appreciated that in example embodiments, the sensor kit 28700 may provide additional types of data to the backend system 28750. For example, the sensor kit 28700 may provide diagnostic data indicating any detected issues (e.g., malfunction, battery levels low, etc.) or potential issues with the sensors 28702 or other devices in the sensor kit 28700.

[1522] In example embodiments, the sensor kit 28700 is configured to self-monitor for failing components (e.g., failing sensors 28702) and to report failing components to the operator. For example, in example embodiments, the edge device 28704 may be configured to detect failure of a sensor 28702 based on a lack of reporting from a sensor, a lack of response to requests (e.g., "pings"), and/or based on unreliable data (e.g., data regularly falling out of the expected sensor readings). In example embodiments, the edge device 28704 can maintain a sensor kit network map indicating where each device in the sensor kit network is located and can provide approximate locations and/or identifiers of failed sensors to a user.

[1523] In example embodiments, the sensor kit 28700 may be implemented to allow post-installation configuration. A post-installation configuration may refer to an update to the sensor kit 28700 by adding devices and/or services to the sensor kit 28700 after the sensor kit 28700 has been installed. In some of these embodiments, users (e.g., operators of the industrial setting 28720) of the system may subscribe to or purchase certain edge "services." For example, the sensor kit 28700 may be configured to execute certain programs installed on one or more devices of the sensor kit 28700 only if the user has a valid subscription or ownership permission to access the edge service supported by the program. When the user no longer has the valid subscription and/or ownership permission, the sensor kit 28700 may preclude execution of those programs. For example, a user may subscribe to unlock AI-based edge services, mesh networking capabilities, self-monitoring services, compression services, in-facility notifications, and the like.

[1524] In example embodiments, users can add new sensors 28702 to the sensor kit post-installation in a plug-and-play-like manner. In some of these embodiments, the edge device 28704 and the sensors 28702 (or other devices to be added to the sensor kit 28700) may include respective short-range communication capabilities (e.g., near-field communication (NFC) chips, RFID chips, BLUETOOTH™ chips, Wi-Fi adapters, and the like). In these embodiments, the sensors 28702 may include persistent storage that stores identifying data (e.g., a sensor identifier value) and any other data that would be used to add the sensor 28702 to the sensor kit 28700 (e.g., an industrial device type, supported communication protocols, and the like). In example embodiments, a user may initiate a post-installation addition to the sensor kit 28700 by pressing a button on the edge device 28704, and/or by bringing the sensor 28702 into the vicinity of the edge device 28704. In example embodiments, in response to a user initiating a post-installation addition to the sensor kit, the edge device 28704 may emit a signal (e.g., a radio frequency). The edge device 28704 may emit the signal, for example, as a result of a human user pushing a button or at a predetermined time interval. The emitted signal may trigger a sensor 28702 proximate enough to receive the signal and to transmit the sensor ID of the sensor 28702 and any other suitable configuration data (e.g., device type, communication protocols, and the like). In response to the sensor 28702 transmitting its configuration data (e.g., sensor ID and other relevant configuration data) to the edge device

28704, the edge device 28704 may add the sensor 28702 to the sensor kit 28702. Adding the sensor 28702 to the sensor kit 28704 may include updating a data store or manifest stored at the edge device 28704 that identifies the devices of the sensor kit 28700 and data relating thereto. Non-limiting examples of data that may be stored in the manifest relating to each respective sensor 28702 may include the communication protocol used by the sensor 28702 to communicate with the edge device 28704 (or intermediate devices), the type of sensor data provided by the sensor 28702 (e.g., vibration sensor data, temperature data, humidity data, etc.), models used to analyze sensor data from the sensor 28702 (e.g., a model identifier), alarm limits associated with the sensor 28702, and the like.

[1525] In example embodiments, the sensor kit 28700 (e.g., the edge device 28704) may be configured to update a distributed ledger 28762 with sensor data captured by the sensor kit 28700. In example embodiments, a distributed ledger 28762 is a Blockchain or any other suitable distributed ledger 28762. The distributed ledger 28762 may be a public ledger or a private ledger. Private ledgers reduce power consumption requirements of maintaining the distributed ledger 28762, while public ledgers consume more power but offer more robust security. In example embodiments, the distributed ledger 28762 may be distributed amongst a plurality of node computing devices 28760. The node computing devices 28760 may be any suitable computing device, including physical servers, virtual servers, personal computing devices, and the like. In example embodiments, the node computing devices 28760 are approved (e.g., via a consensus mechanism) before the node computing devices 28760 may participate in the distributed ledger. In example embodiments, the distributed ledger 28762 may be privately stored. For example, a distributed ledger may be stored amongst a set of preapproved node computing devices, such that the distributed ledger 28762 is not accessible by non-approved devices. In example embodiments, the node computing devices 28760 are edge devices 28704 of the sensor kit 28702 and other sensor kits 28702.

[1526] In example embodiments, the distributed ledger 28762 is comprised of a set of linked data structures (e.g., blocks, data records, etc.), such that the linked data structures form an acyclic graph. For purposes of explanation, the data structures will be referred to as blocks. In example embodiments, each block may include a header that includes a unique ID of the block and a body that includes the data that is stored in the block, and a pointer. In example embodiments, the pointer is the block ID of a parent block of the block, wherein the parent block is a block that was created prior to the block being written. The data stored in a respective block can be sensor data captured by a respective sensor kit 28700. Depending on the implementation, the types of sensor data and the amount of sensor data stored in a respective body of a block may vary. For example, a block may store a set of sensor measurements from one or more types of sensors 28702 of the sensor kit 28700 captured over a period of time (e.g., sensor data 28702 captured from all of the sensors 28702 in the sensor kit 28700 over a period one hour or one day) and metadata relating thereto (e.g., sensor identifiers of each sensor measurement and a timestamp of each sensor measurement or group of sensor measurements). In example embodiments, a block may store sensor measurements determined to be anomalous (e.g., outside a standard deviation of expected sensor

measurements or deltas in sensor measurements that are above a threshold) and/or sensor measurements indicative of an issue or potential issue, and related metadata (e.g., sensor IDs of each sensor measurement and a timestamp of each sensor measurement or group of sensor measurements). In example embodiments, the sensor data stored in a block may be compressed and/or encoded sensor data, such that the edge device 28704 compresses/encodes the sensor data into a more compact format. In example embodiments, the edge device 28704 may generate a hash of the body, such that the contents of the body (e.g., block ID of the parent block and the sensor data) are hashed and cannot be altered without changing the value of the hash. In example embodiments, the edge device 28704 may encrypt the content within the block, so that the content may not be read by unauthorized devices.

[1527] As mentioned, the distributed ledger 28762 may be used for different purposes. In example embodiments, the distributed ledger 28762 may further include one or more smart contracts. A smart contract is a self-executing digital contract. A smart contract may include code (e.g., executable instructions) that defines one or more conditions that trigger one or more actions. A smart contract may be written by a developer in a scripting language (e.g., JavaScript), an object code language (e.g., Java), or a compiled language (e.g., C++ or C). Once written, a smart contract may be encoded in a block and deployed to the distributed ledger 28762. In example embodiments, the backend system 28750 is configured to receive the smart contract from a user and write the smart contract to a respective distributed ledger 28762. In example embodiments, an address of the smart contract (e.g., the block ID of the block containing the smart contract) may be provided to one or more parties to the smart contract, such that respective parties may invoke the smart contract using the address. In example embodiments, the smart contract may include an API that allows a party to provide data (e.g., addresses of blocks) and/or to transmit data (e.g., instructions to transfer funds to an account).

[1528] In example implementations, an insurer may allow insured owners and/or operators of an industrial setting 28720 to agree to share sensor data with the insurer to demonstrate that the equipment in the facility is functioning properly and, in return, the insurer may issue a rebate or refund to the owners and/or operators if the owners and/or operators are compliant with an agreement with the insurers. Compliance with the agreement may be verified electronically by participant nodes in the distributed ledger and/or the sensor kit 28700 via a smart contract. In example embodiments, the insurer may deploy the smart contract (e.g., by adding the smart contract to a distributed ledger 28762) that triggers the issuance of rebates or refunds on portions of insurance premiums when the sensor kit 28700 provides sufficient sensor data to the insurer via the distributed ledger that indicates the facility is operating without issue. In some of these embodiments, the smart contract may include a first condition that requires a certain amount of sensor data to be reported by a facility and a second condition that each instance of the sensor data equals a value (e.g., there are no classified or predicted issues) or range of values (e.g., all sensor measurements are within a predefined range of values). In example embodiments, the action taken in response to one or more of the conditions being met may be to deposit funds (e.g., a wire transfer or cryptocurrency) into an account. In this example, the edge device 28704 may write blocks

containing sensor data to the distributed ledger. The edge device 28704 may also provide the addresses of these blocks to the smart contract (e.g., using an API of the smart contract). Upon the smart contract verifying the first and second conditions of the contract, the smart contract may initiate the transfer of funds from an account of the insurer to the account of the insured.

[1529] In another example, a regulatory body (e.g., a state, local, or federal regulatory agency) may require facility operators to report sensor data to ensure compliance with one or more regulations. For instance, the regulatory body may regulate food inspection facilities, pharmaceutical manufacturing facilities, e.g., manufacturing facility 1700, indoor agricultural facilities, e.g., indoor agricultural facility 1800, offshore oil extraction facilities, e.g., underwater industrial facility 1900, or the like. In example embodiments, the regulatory body may deploy a smart contract that is configured to receive and verify the sensor data from an industrial setting 28720, and in response to verifying the sensor data issues a compliance token (or certificate) to an account of the facility owner. In some of these embodiments, the smart contract may include a condition that requires a certain amount of sensor data to be reported by a facility and a second condition that requires the sensor data to be compliant with the reporting regulations. In this example, the edge device 28704 may write blocks containing sensor data to the distributed ledger 28762. The edge device 28704 may also provide the addresses of these blocks to the smart contract (e.g., using an API of the smart contract). Upon the smart contract verifying the first and second conditions of the contract, the smart contract may generate a token indicating compliance by the facility operator and may initiate the transfer of funds to an account (e.g., a digital wallet) associated with the facility.

[1530] A distributed ledger 28762 may be adapted for additional or alternative applications without departing from the scope of the disclosure.

[1531] FIGS. 139, 140, and 141 illustrate example configurations of a sensor kit network 28800. Depending on the sensor kit 28700 and the industrial setting 28720 that the sensor kit 28700 is installed in, the sensor kit network 28800 may communicate in different manners.

[1532] FIG. 139 illustrates an example sensor kit network 28800A that is a star network. In these embodiments, the sensors 28702 communicate directly with the edge device 28704. In these embodiments, the communication protocol(s) utilized by the sensor devices 28702 and the edge device 28704 to communicate are based on one or more of the physical area of the sensor kit network 28702, the power sources available, and the types of sensors 28702 in the sensor kit 28700. For example, in settings where the area being monitored is a relatively small area and where the sensors 28702 are not able to connect to a power supply, the sensors 28702 may be fabricated with a BLUETOOTH LOW ENERGY™ (BLE) microchip that communicates using a BLUETOOTH LOW ENERGY™ protocol (e.g., the BLUETOOTH™ 5 protocol maintained by the Bluetooth Special Interest Group). In another example, in a relatively small area where lots of sensors 28702 are to be deployed, the sensors 28702 may be fabricated with the Wi-Fi microchip that communicates using the IEEE 802.11 protocol. In the embodiments of FIG. 139, the sensors 28702 may be configured to perform one-way or two-way communication. In embodiments where the edge device 28704 does not need to communicate data and/or instructions to the sensors 28702,

the sensors 28702 may be configured for one-way communication. In embodiments where the edge device 28704 does communicate data and/or instructions to the sensors 28702, the sensors 28702 may be configured with transceivers that perform two-way communication. A star network may be configured with devices having other suitable communication devices without departing from the scope of the disclosure.

[1533] FIG. 140 illustrates an example sensor kit network 28800B that is a mesh network where the nodes (e.g., sensors 28702) connect to each other directly, dynamically, and/or non-hierarchically to cooperate with one another to efficiently route data to and from the edge device 28704. In example embodiments, the devices in the mesh network (e.g., the sensors 28702, the edge device 28704, and/or any other devices in the sensor kit network 28800B) may be configured to self-organize and self-configure the mesh network, such that the sensors 28702 and/or the edge device 28704 may determine which devices route data on behalf of other devices, and/or redundancies for transmission should a routing node (e.g., sensor 28702) fail. In example embodiments, the sensor kit 28700 may be configured to implement a mesh network in industrial settings 28720 where the area being monitored is relatively large (e.g., greater than 28700 meters in radius from the edge device 28704) and/or where the sensors 28702 in the sensor kit 28700 are intended to be installed in close proximity to one another. In the latter scenario, the power consumption of each individual sensor 28702 may be reduced in comparison to sensors 28702 in a star network, as the distance that each respective sensor 28702 needs to transmit over is relatively less than the distance that the respective sensor 28702 would need to transmit over in a star network. In example embodiments, a sensor 28702 may be fabricated with a ZIGBEE® microchips, a Digi XBee® microchip, a BLUETOOTH LOW ENERGY™ microchip, and/or any other suitable communication devices configured to participate in a mesh network.

[1534] FIG. 141 illustrates an example of a sensor kit network 28800C that is a hierarchical network. In these embodiments, the sensor kit 28700 includes a set of collection devices 28806. A collection device 28806 may refer to a non-sensor device that receives sensor data from a sensor device 28704 and routes the sensor data to an edge device 28704, either directly or via another collection device 28806. In example embodiments, a hierarchical network may refer to a network topography where one or more intermediate devices (e.g., collection devices 28806) route data from one or more respective peripheral devices (e.g., sensor devices 28702) to a central device (e.g., edge device 28704). A hierarchical network may include wired and/or wireless connections. In example embodiments, a sensor device 28702 may be configured to communicate with a collection device 28806 via any suitable communication device (e.g., BLUETOOTH LOW ENERGY™ microchips, Wi-Fi microchips, ZIGBEE® microchips, or the like). In example embodiments, hierarchical sensor kit networks may be implemented in industrial settings 28720 where power sources are available to power the collection devices 28806 and/or where the sensors 28702 are likely to be spaced too far apart to support a reliable mesh network.

[1535] The examples of FIGS. 139-141 are provided for examples of different topologies of a sensor kit network. These examples are not intended to limit the types of sensor kit networks 28800 that may be formed by a sensor kit 28700. Furthermore, sensor kit networks 28800 may be

configured as hybrids of star networks, hierarchical networks, and/or mesh networks, depending on the industrial settings 28720 in which respective sensor kits 28800 are being deployed.

[1536] FIG. 289A illustrates an example IoT sensor 28702 (or sensor) according to embodiments of the present disclosure. Embodiments of the IoT sensor 28702 may include, but are not limited to, one or more sensing components 28902, one or more storage devices 28904, one or more power supplies 28906, one or more communication devices 28908, and a processing device 28910. In example embodiments, the processing device 28910 may execute an edge reporting module 28912.

[1537] A sensor 28702 includes at least one sensing component 28902. A sensing component 28902 may be any digital, analog, chemical, and/or mechanical component that outputs raw sensor data to the processing device 28910. It is appreciated that different types of sensors 28702 are fabricated with different types of sensing components. In example embodiments, sensing components 28902 of an inertial sensor may include one or more accelerometers and/or one or more gyroscopes. In example embodiments, sensing components 28902 of a temperature sensor may include one or more thermistors or other temperature sensing mechanisms. In example embodiments, sensing components 28902 of a heat flux sensor may include, for example, thin film sensors, surface mount sensors, polymer-based sensors, chemical sensors and others. In example embodiments, sensing components 28902 of a motion sensor may include a LIDAR device, a radar device, a sonar device, or the like. In example embodiments, sensing components 28902 of an occupancy sensor may include a surface being monitored for occupancy, a pressure activated switch embedded under the surface of the occupancy sensor and/or a piezoelectric element integrated into the surface of the occupancy sensor, such that an electrical signal is generated when an object occupies the surface being monitored for occupancy. In example embodiments, sensing components 28902 of a humidity sensor may include a capacitive element (e.g., a metal oxide between to electrodes) that outputs an electrical capacity value corresponding to the ambient humidity; a resistive element that includes a salt medium having electrodes on two sides of the medium, whereby the variable resistance measured at the electrodes corresponds to the ambient humidity; and/or a thermal element that includes a first thermal sensor that outputs a temperature of a dry medium (e.g., dry nitrogen) and a second thermal sensor that outputs an ambient temperature of the sensor's environment, such that the humidity is determined based on the change, i.e., the delta, between the temperature in the dry medium and the ambient temperature. In example embodiments, sensing components 28902 of a vibration sensor may include accelerometer components, position sensing components, torque sensing components, and others. It is appreciated that the list of sensor types and sensing components thereof is provided for example. Additional or alternative types of sensors and sensing components may be integrated into a sensor 28702 without departing from the scope of the disclosure. Furthermore, in example embodiments, the sensors 28702 of a sensor kit 28700 may include audio, visual, or audio/visual sensors, in addition to non-audio/visual sensors 28702 (i.e., sensors that do not capture video or audio). In these embodiments, the sensing components 28992 may include a camera and/or one or more microphones. In example embodiments, the microphones may be directional microphones, such that a direction of a source of audio may be determined.

[1538] A storage device 28904 may be any suitable medium for storing data that is to be transmitted to the edge device 28704. In example embodiments, a storage device 28904 may be a persistent storage medium, such as a flash memory device. In example embodiments, a storage device 28904 may be a transitory storage medium, such as a random access memory device. In example embodiments, a storage device 28904 may be a circuit configured to store charges, whereby the magnitude of the charge stored by the component is indicative of a sensed value, or incremental counts. In these embodiments, this type of storage device 28904 may be used where power availability and size are concerns, and/or where the sensor data is count-based (e.g., a number of detection events). It is appreciated that any other suitable storage devices 28904 may be used. In example embodiments, the storage device 28904 may include a cache 28914, such that the cache 28914 stores sensor data that is not yet reported to the edge device 28704. In these embodiments, the edge reporting module 28912 may clear the cache 28914 after the sensor data being stored in the cache 28914 is transmitted to the edge device 28704.

[1539] A power supply 28906 is any suitable component that provides power to the other components of the sensor 28702, including the sensing components 28902, storage devices 28904, communication devices 28906, and/or the processing device 28908. In example embodiments, a power supply 28906 includes a wired connection to an external power supply (e.g., alternating current delivered from a power outlet, or direct current delivered from a battery or solar power supply). In example embodiments, the power supply 28906 may include a power inverter that converts alternating currents to direct currents (or vice-versa). In example embodiments, a power supply 28906 may include an integrated power source, such as a rechargeable lithium ion battery or a solar element. In example embodiments, a power supply 28906 may include a self-powering element, such as a piezoelectric element. In these embodiments, the piezoelectric element may output a voltage upon a sufficient mechanical stress or force being applied to the element. This voltage may be stored in a capacitor and/or may power a sensing element 28902. In example embodiments, the power supply may include an antenna (e.g., a receiver or transceiver) that receives a radio frequency that energizes the sensor 28702. In these embodiments, the radio frequency may cause the sensor 28702 to "wake up" and may trigger an action by the sensor 28702, such as taking sensor measurements and/or reporting sensor data to the edge device 28704. A power supply 28906 may include additional or alternative components as well.

[1540] In example embodiments, a communication device 28908 is a device that enables wired or wireless communication with another device in the sensor kit network 28800. In most sensor kit configurations 28700, the sensors 28702 are configured to communicate wirelessly. In these embodiments, a communication device 28908 may include a transmitter or transceiver that transmits data to other devices in the sensor kit network 28800. Furthermore, in some of these embodiments, communication devices 28908 having transceivers may receive data from other devices in the sensor kit network 200. In wireless embodiments, the transceiver may be integrated into a chip that is configured to perform communication using a respective communication protocol. In example embodiments, a communication device 28908 may be a ZIGBEE® microchip, a Digi XBee® microchip, a BLUETOOTH™ microchip, a BLUETOOTH LOW

ENERGY™ microchip, a Wi-Fi microchip, or any other suitable short-range communication microchip. In embodiments where the sensor kit 200 supports a mesh network, the communication device 28908 may be a microchip that implements a communication protocol that supports mesh networking (e.g., ZIGBEE PRO mesh networking protocol, BLUETOOTH™ Mesh, 802.11a/b/g/n/ac, and the like). In these embodiments, a communication device 28908 may be configured to establish the mesh network and handle the routing of data packets received from other devices in accordance with the communication protocol implemented by the communication device 28908. In example embodiments, a sensor 28702 may be configured with two or more communication devices 28908. In these embodiments, the sensors 28702 may be added to different sensor kit 28700 configurations and/or may allow for flexible configuration of the sensor kit 28702 depending on the industrial setting 28720.

[1541] In example embodiments, the processing device 28910 may be a microprocessor. The microprocessor may include memory (e.g., read-only memory (ROM)) that stores computer-executable instructions and one or more processors that execute the computer-executable instructions. In example embodiments, the processing device 28910 executes an edge reporting module 28912. In example embodiments, the edge reporting module 28912 is configured to transmit data to the edge device 28704. Depending on the configuration of the sensor kit network 28800 and location of the sensors 28702 with respect to the edge device 28704, the edge reporting module 28912 may transmit data (e.g., sensor data) either directly to the edge device 28704, or to an intermediate device (e.g., a collection device 206 or another sensor device 28702) that routes the data towards the edge device 28704. In example embodiments, the edge reporting module 28912 obtains raw sensor data from a sensing component 28902 or from a storage device 28904 and packetizes the raw sensor data into a reporting packet 28920.

[1542] FIG. 289B illustrates an example reporting packet 28920 according to some embodiments of the present disclosure. In some of these embodiments, the edge reporting module 28912 may populate a reporting packet template to obtain a reporting packet 28920. In example embodiments, a reporting packet 28920 may include a first field 28922 indicating a sensor ID of the sensor 28702 and a second field 28926 indicating the sensor data. Additionally, the reporting packet 28920 may include additional fields, such as a routing data field 28924 indicating a destination of the packet (e.g., an address or identifier of the edge device 28704), a time stamp field 28928 indicating a time stamp, and/or a checksum field 28930 indicating a checksum (e.g., a hash value of the contents of the reporting packet). The reporting packet may include additional or alternative fields (e.g., error codes) without departing from the scope of the disclosure.

[1543] Referring back to FIG. 142, in example embodiments, the edge reporting module 28912 may generate a reporting packet 28920 for each instance of sensor data. Alternatively, the edge reporting module 28912 may generate a reporting packet 28920 that includes a batch of sensor data (e.g., the previous N sensor readings or all the sensor readings maintained in a cache 28914 of the sensor 28702 since the cache 28914 was last purged). Upon generating a reporting packet 28920, the edge reporting module 28912 may output the reporting packet 28920 to the communication device 28908, which transmits the reporting packet 28920 to the edge device 28704 (either directly

or via one or more intermediate devices). The edge reporting module 28912 may generate and transmit reporting packets 28920 at predetermined intervals (e.g., every second, every minute, every hour), continuously, or upon being triggered (e.g., upon being activated via the power supply or upon being command by the edge device 28704).

[1544] In example embodiments, the edge reporting module 28912 instructs the sensing component(s) 28902 to capture sensor data. In example embodiments, the edge reporting module 28912 may instruct a sensing component 28902 to capture sensor data at predetermined intervals. For example, the edge reporting module 28912 may instruct the sensing component 28902 to capture sensor data every second, every minute, or every hour. In example embodiments, the edge reporting module 28912 may instruct a sensing component 28902 to capture sensor data upon the power supply 28906 being energized. For example, the power supply 28906 may be energized by a radio frequency or upon a pressure-switch being activated and closing a circuit. In example embodiments, the edge reporting module 28912 may instruct a sensing component 28902 to capture sensor data in response to receiving a command to report sensor data from the edge device 28704 or a human user (e.g., in response to the user pressing a button).

[1545] In example embodiments, a sensor 28702 includes a housing (not shown). The sensor housing may have any suitable form factor. In embodiments where the sensor 28702 is being used outdoors, the sensor may have a housing that is waterproof and/or resistant to extreme cold and/or extreme heat. In example embodiments, the housing may have suitable coupling mechanisms to removably couple to an industrial component.

[1546] The foregoing is an example of a sensor 28702. The sensor 28702 may have additional or alternative components without departing from the scope of the disclosure.

[1547] FIG. 290 illustrates an example of an edge device 28704. In example embodiments, the edge device 28704 may include a storage system 29002, a communication system 29004, and a processing system 29006. The edge device 28704 may include additional components not shown, such as a power supply, a user interface, and the like.

[1548] The storage system 29002 includes one or more storage devices. The storage devices may include persistent storage mediums (e.g., flash memory drive, hard disk drive) and/or transient storage devices (e.g., RAM). The storage system 29002 may store one or more data stores. A data store may include one or more databases, tables, indexes, records, filesystems, folders and/or files. In the illustrated embodiments, the storage device stores a configuration data store 29010, a sensor data store 29012, and a model data store 29014. A storage system 29002 may store additional or alternative data stores without departing from the scope of the disclosure.

[1549] In example embodiments, the configuration data store 29010 stores data relating to the configuration of the sensor kit 28700, including the devices of the sensor kit 28700. In example embodiments, the configuration data store 29010 may maintain a set of device records. The device records may indicate a device identifier that uniquely identifies a device of the sensor kit 28700. The device records may further indicate the type of device (e.g., a sensor, a collection device, a gateway device, etc.). In embodiments where the network paths from each device to the edge device 28704 do not change, a device record may also indicate the network path of the device to

the edge device 28704 (e.g., any intermediate devices in the device's network path). In the case that a device record corresponds to a sensor 28702, the device record may indicate the type of sensor (e.g., a sensor type identifier) and/or a type of data that is provided by the sensor 28702.

[1550] In example embodiments, the configuration data store 29010 may maintain a set of sensor type records, where each record corresponds to a different type of sensor 28702 in the sensor kit 28700. A sensor type record may indicate a type identifier that identifies the type of sensor and/or the type of sensor data provided by the sensor. In example embodiments, a sensor type record may further indicate relevant information relating to the sensor data, including maximum or minimum values of the sensor data, error codes output by sensors 28702 of the sensor type, and the like.

[1551] In example embodiments, the configuration data store 29010 may maintain a map of the sensor kit network 200. The map of the sensor kit network 200 may indicate a network topology of the sensor kit network 200, including network paths of the collection of devices in the sensor kit 28700. In example embodiments, the map may include physical locations of the sensors as well. The physical location of a sensor 28702 may be defined as a room or area that the sensor 28702 is in, a specific industrial component that the sensor 28702 is monitoring, a set of coordinates relative of the edge device 28704 (e.g., x, y, z coordinates relative to the edge device 28704, or an angle and distance of the sensor 28702 relative to the edge device 28704), an estimated longitude and latitude of the sensor 28702, or any other suitable format of relative or absolute location determination and/or measurement.

[1552] In example embodiments, a sensor data store stores 29012 stores sensor data collected from the sensors 28702 of the sensor kit 28700. In example embodiments, the sensor data store 29012 maintains sensor data that is collected over a period of time. In some of these embodiments, the sensor data store 29012 may be a cache that stores sensor data until it is reported and backed up at the backend system 28750. In these embodiments, the cache may be cleared when sensor data is reported to the backend system 28750. In example embodiments, the sensor data store 29012 stores all sensor data collected by the sensor kit 29012. In these embodiments, the sensor data store 29012 may provide a backup for all the sensor data collected by the sensor kit 28700 over time, thereby ensuring that the owner of the sensor kit 28700 maintains ownership of its data.

[1553] In example embodiments, a model data store 29014 stores machine-learned models. The machine-learned models may include any suitable type of models, including neural networks, deep neural networks, recursive neural networks, Bayesian neural networks, regression-based models, decision trees, prediction trees, classification trees, Hidden Markov Models, and/or any other suitable types of models. A machine-learned model may be trained on training data, which may be expert generated data, historical data, and/or outcome-based data. Outcome-based data may be data that is collected after a prediction or classification is made that indicates whether the prediction or classification was correct or incorrect and/or a realized outcome. A training data instance may refer to a unit of training data that includes a set of features and a label. In example embodiments, the label in a training data instance may indicate a condition of an industrial component or an industrial setting 28720 at a given time. Examples of conditions will vary greatly depending on the industrial setting 28720 and the conditions that the machine-learned model is being trained to predict or

classify. Examples of labels in a manufacturing facility may include, but are not limited to, no issues detected, a mechanical failure of a component, an electrical failure of a component, a chemical leak detected, and the like. Examples of labels in a mining facility may include, but are not limited to, no issues detected, an oxygen deficiency, the presence of a toxic gas, a failing structural component, and the like. Examples of labels in an oil and/or gas facility (e.g., oil field, gas field, oil refinery, pipeline) may include, but are not limited to, no issues detected, a mechanical failure of a component (e.g., a failed valve or failed O-ring), a leak, and the like. Examples of labels in an indoor agricultural facility may include, but are not limited to, no issues detected, a plant died, a plant wilted, a plant turned a certain color (e.g., brown, purple, orange, or yellow), mold found, and the like. In each of these examples, there are certain features that may be relevant to a condition and some features that may have little or no bearing on the condition. Through a machine-learning process (which may be performed at the backend system 28750 or another system), the model is trained to determine predictions or classifications based on a set of features. Thus, the set of features in a training data instance may include sensor data that is temporally proximate to a time when a condition of the industrial component or industrial setting 28720 occurred (e.g., the label associated with the industrial component or industrial setting 28720).

[1554] In example embodiments, the machine-learned models may include prediction models that are used to predict potential issues relating to an industrial component being monitored. In some of these embodiments, a machine-learned model may be trained on training data (expert generated data and/or historical data) that corresponds to one or more conditions relating to a particular component. In some of these embodiments, the training data sets may include sensor data corresponding to scenarios where maintenance or some intervening action was later required and sensor data corresponding to scenarios where maintenance or some intervening action was ultimately not required. In these example embodiments, the machine-learned model may be used to determine a prediction of one or more potential issues that may arise with respect to one or more industrial components being monitored and/or the industrial setting 28720 being monitored.

[1555] In example embodiments, the machine-learned models may include classification models that classify a condition of an industrial component being monitored and/or the industrial setting 28720. In some of these embodiments, a machine-learned model may be trained on training data (e.g., expert generated data and/or historical data) that corresponds to one or more conditions relating to a particular component. In some of these embodiments, the training data sets may include sensor data corresponding to scenarios where respective industrial components and/or respective industrial settings 28720 were operating in a normal condition and sensor data where the respective industrial components and/or respective industrial settings 28720 were operating in an abnormal condition. In training data instances where there was an abnormal condition, the training data instance may include a label indicating the type of abnormal condition. For example, a training data instance corresponding to an indoor agricultural facility that was deemed too humid for ideal growing conditions may include a label that indicates the facility was too humid.

[1556] In example embodiments, the communication system 29004 includes two or more communication devices, including at least one internal communication device that communicates

with the sensor kit network 200 and at least one external communication device that communicates with a public communication network (e.g., the Internet) either directly or via a gateway device. The at least one internal communication devices may include BLUETOOTH™ chips, Zigbee chips, XBee chips, Wi-Fi chips, and the like. The selection of the internal communication devices may depend on the environment of the industrial setting 28720 and the impacts thereof on the sensors 28702 to be installed therein (e.g., whether the sensors 28702 have reliable power sources, whether the sensors 28702 will be spaced in proximity to one another, whether the sensors 28702 need to transmit through walls, and the like). The external communication devices may perform wired or wireless communication. In example embodiments, the external communication devices may include cellular chipsets (e.g., 4G or 5G chipsets), Ethernet cards, satellite communication cards, or other suitable communication devices. The external communication device(s) of an edge device 28704 may be selected based on the environment of the industrial setting 28720 (e.g., indoors v. outdoors, thick walls that prevent wireless communication v. thin walls that allow wireless communication, located near cellphone towers v. located in remote areas) and the preferences of an operator of the industrial setting 28720 (e.g., the operator allows the edge device 28704 to access a private network of the industrial setting 28720, or the operator does not allow the edge device 28704 to access a private network of the industrial setting 28720).

[1557] In example embodiments, the processing system 29006 may include one or more memory devices (e.g., ROM and/or RAM) that store computer-executable instructions and one or more processors that execute the computer-executable instructions. The processing system 29006 may execute one or more of a data processing module 29020, an encoding module 29022, a quick-decision AI module 29024, a notification module 29026, a configuration module 29028, and a distributed ledger module 29030. The processing system 29006 may execute additional or alternative modules without departing from the scope of the disclosure. Furthermore, the modules discussed herein may include submodules that perform one or more functions of a respective module.

[1558] In example embodiments, the data processing module 29020 receives sensor data from the sensor kit network 200 and performs one or more data processing operations on the received sensor data. In example embodiments, the data processing module 29020 receives reporting packets 320 containing sensor data. In some of these embodiments, the data processing module 29020 may filter data records that are duplicative (e.g., filtering out one out of two reporting packets 320 received from two respective sensors monitoring the same component for redundancy). The data processing module 29020 may additionally or alternatively filter and/or flag reporting packets 320 containing sensor data that is clearly erroneous (e.g., sensor not within a tolerance range given the type of sensor 28702 or contains an error code). In example embodiments, the data processing module 29020 may store and/or index the sensor data in the sensor data store.

[1559] In example embodiments, the data processing module 29020 may aggregate sensor data received over a period of time from the sensors 28702 of the sensor kit 28700 or a subset thereof and may transmit the sensor data to the backend system 28750. In transmitting sensor data to the backend system 28750, the data processing module 29020 may generate a sensor kit reporting

packet that includes one or more instances of sensor data. The sensor data in the sensor kit reporting packet may be compressed or uncompressed. In example embodiments, the sensor kit reporting packet may indicate a sensor kit identifier that identifies the source of the data packet to the backend system 28750. In example embodiments, the data processing module 29020 may transmit the sensor data upon receipt of the sensor data from a sensor 28702, at predetermined intervals (e.g., every second, every minute, every hour, every day), or in response to a triggering condition (e.g., a prediction or classification that there is an issue with an industrial component or the industrial setting 28720 based on received sensor data). In example embodiments, the sensor data may be encoded/compressed, such that sensor data collected from multiple sensors 28702 and/or over a period of time may be more efficiently transmitted. In example embodiments, the data processing module 29020 may leverage the quick-decision AI module 29024 to determine whether the industrial components of the industrial setting 28720 and/or the industrial setting 28720 itself is likely in a normal condition. If the quick-decision AI module 29024 determines that the industrial components and/or the industrial setting 28720 are in a normal condition with a high degree of certainty, then the data processing module 29020 may delay or forgo transmitting the sensor data used to make the classification to the backend system 28750. Additionally or alternatively, if the quick-decision AI module 29024 determines that the industrial components and/or the industrial setting 28720 are in a normal condition with a high degree of certainty, then the data processing module 29020 may compress the sensor data and may be compressed at a greater rate. The data processing module 29020 may perform additional or alternative functions without departing from the scope of the disclosure.

[1560] In example embodiments, the encoding module 29022 receives sensor data and may encode, compress, and/or encrypt the sensor data. The encoding module 29022 may employ other techniques to compress the sensor data. In example embodiments, the encoding module 29022 may employ horizontal or compression techniques to compress the sensor data. For example, the encoding module 29022 may use the Lempel-Zev-Welch algorithm or variations thereof. In example embodiments, the encoding module 522 may represent sensor data in an original integer or "counts format" and with relevant calibration coefficients and offsets at the time of collection. In these embodiments, the coefficients and offsets may be coalesced at the time of collection when a precise signal path is known, such that one floating-point coefficient and one integer offset is stored for each channel.

[1561] In example embodiments, the encoding module 29022 may employ one or more codecs to compress the sensor data. The codecs may be proprietary codecs and/or publicly available codecs. In example embodiments, the encoding module 29022 may use a media compression codec (e.g., a video compression codec) to compress the sensor data. For example, the encoding module 29022 may normalize the sensor data into values that fall within a range and format of a media frame (e.g., normalizing sensor data into acceptable pixel values for inclusion into a video frame) and may embed the normalized sensor data into the media frame. The encoding module 29022 may embed the normalized sensor data collected from the sensors 28702 of the sensor kit 28700 into the media frame according to a predefined mapping (e.g., a mapping of respective sensors 28702

to one or more respective pixels in a media frame). The encoding module 29022 may generate a set of consecutive media frames in this manner and may compress the media frames using a media codec (an H.264/MPEG-4 codec, an H.265/MPEG-H codec, an H.263/MPEG-4 codec, proprietary codecs, and the like) to obtain a sensor data encoding. The encoding module 29022 may then transmit sensor data encoding to the backend system, which may decompress and recalculate the sensor data based on the normalized values. In these embodiments, the codec used for compression and the mappings of sensors to pixels may be selected to reduce lossiness or to increase compression rates. Furthermore, the foregoing technique may be applied to sensor data that tends to be more static and less changing between samplings and/or where sensor data collected from different sensors tend to have little variation when sampled at the same time. The encoding module 29022 may employ additional or alternative encoding/compression techniques without departing from the scope of the disclosure.

[1562] In example embodiments, the quick-decision AI module 29024 may utilize a limited set of machine-learned models to generate predictions and/or classifications of a condition of an industrial component being monitored and/or of the industrial setting 28720 being monitored. In example embodiments, the quick-decision AI module 29024 may receive a set of features (e.g., one or more sensor data values) and request for a specific type of prediction or classification based thereon. In example embodiments, the quick-decision AI module 29024 may leverage a machine-learned model corresponding to the requested prediction or classification. The quick-decision AI module 29024 may generate a feature vector based on the received features, such that the feature vector includes one or more sensor data values obtained from one or more sensors 28702 of the sensor kit 28700. The quick-decision AI module 29024 may feed the feature vector to the machine-learned model. The machine-learned model may output a prediction or classification and a degree of confidence in the prediction or classification. In example embodiments, the quick-decision AI module 29024 may output the prediction or classification to the data processing module 29020 (or another module that requested a prediction or classification). For example, in embodiments the data processing module 29020 may use classifications that the industrial components and/or the industrial setting 28720 are in a normal condition to delay or forgo transmission of sensor data and/or to compress sensor data. In example embodiments, the data processing module 29020 may use a prediction or classification that the industrial components and/or the industrial setting 28720 are likely to encounter a malfunction to transmit uncompressed sensor data to the backend system 28750, which may further analyze the sensor data and/or notify a human user of a potential issue.

[1563] In example embodiments, the notification module 29026 may provide notifications or alarms to users based on the sensor data. In some of these embodiments, the notification module 29026 may apply a set of rules that trigger a notification or alarm if certain conditions are met. The conditions may define sensor data values that are strongly correlated with an undesirable (e.g., emergency) condition. Upon receiving sensor data from the data processing module 29020, the notification module 29026 may apply one or more rules to the sensor data. If the conditions to trigger an alarm or notification are met, the notification module 29026 may issue an alarm or notification to a human user. The manner by which an alarm or notification is provided to the

human user (e.g., to a user device, or triggering an audible alarm) may be predefined or, in example embodiments, may be defined by an operator of the industrial setting 28720.

[1564] In example embodiments, the configuration module 29028 configures the sensor kit network 200. In example embodiments, the configuration module 29028 may transmit configuration requests to the other devices in the sensor kit 28700, upon the sensors 28702, edge device 28704, and any other devices being installed in the industrial setting 28720. In some of these embodiments, the sensors 28702 and/or other devices may establish a mesh network or a hierarchical network in response to the configuration requests. In example embodiments, the sensors 28702 and other devices in the sensor kit network may respond to the configuration requests, in response to the configuration requests. In example embodiments, the configuration module 29028 may generate device records corresponding to the devices that responded based on the device IDs of those devices and any additional data provided in the responses to the configuration requests.

[1565] In example embodiments, the configuration module 29028 adds new devices to the sensor kit 28700. In these embodiments, the configuration module 29028 adds new sensors 28702 to the sensor kit 28700 post-installation in a plug-and-play-like manner. In some of these embodiments, the communication devices 29004, 308 of the edge device 28704 and the sensors 28702 (or other devices to be added to the sensor kit 28700) may include respective short-range communication capabilities (e.g., near-field communication (NFC) chips). In these embodiments, the sensors 28702 may include persistent storage that stores identifying data (e.g., a sensor id value) and any other data that would be used to add the sensor to the sensor kit (e.g., device type, supported communication protocols, and the like). In response to a user initiating a post-installation addition to the sensor kit 28700 (e.g., the user pressing a button on the edge device 28704 and/or bringing the sensor 28702 into the vicinity of the edge device 28704), the configuration module 29028 may cause the communication system 29004 to emit a signal (e.g., a radio frequency). The emitted signal may trigger a sensor 28702 proximate enough to receive the signal to transmit its sensor ID and any other suitable configuration data (e.g., device type, communication protocols, and the like). In response to the sensor 28702 transmitting its configuration data (sensor ID and other relevant configuration data) to the edge device 28704, the configuration module 29028 may add the new sensor 28702 to the sensor kit 28702. In example embodiments, adding the sensor 28702 to the sensor kit 28704 may include generating a new device record corresponding to the new sensor 28702 based on the sensor id updating the configuration data store 29010 with the new device record. The configuration module 29028 may add a new sensor 28702 to the sensor kit 28700 in any other suitable manner.

[1566] In example embodiments, the edge device 28704 may include a distributed ledger module 29030. In example embodiments, the distributed ledger module 29030 may be configured to update a distributed ledger 28762 with sensor data captured by the sensor kit 28700. In example embodiments, the distributed ledger may be distributed amongst a plurality of node computing devices 28760. As discussed, in example embodiments, a distributed ledger 28762 is comprised of

a set of linked data structures (e.g., blocks, data records, etc.). For purposes of explanation, the data structures will be referred to as blocks.

[1567] As discussed, each block may include a header that includes a unique ID of the block and a body that includes the data that is stored in the block and a pointer of a parent block. In example embodiments, the pointer in the block is the block ID of a parent block of the block. The data stored in a respective block can be sensor data captured by a respective sensor kit 28700. Depending on the implementation, the types of sensor data and the amount of sensor data stored in a respective body of a block may vary. For example, a block may store a set of sensor measurements from one or more types of sensors 28702 in the sensor kit 28700 captured over a period of time (e.g., sensor data 28702 captured from all of the sensors 28702 in the sensor kit 28700 over a period one hour or one day) and metadata relating thereto (e.g., sensor IDs of each sensor measurement and a timestamp of each sensor measurement or group of sensor measurements). In example embodiments, a block may store sensor measurements determined to be anomalous (e.g., outside a standard deviation of expected sensor measurements or deltas in sensor measurements that are above a threshold) and/or sensor measurements indicative of an issue or potential issue, and related metadata (e.g., sensor IDs of each sensor measurement and a timestamp of each sensor measurement or group of sensor measurements). In example embodiments, the sensor data stored in a block may be compressed and/or encoded sensor data, such that the encoding module 29022 compresses/encodes the sensor data into a more compact format. In example embodiments, the distributed ledger module 29030 may generate a hash of the body, such that the contents of the body (e.g., block ID of the parent block and the sensor data) are hashed and cannot be altered without changing the value of the hash. In example embodiments, the distributed ledger module 29030 may encrypt the content within the block, so that the content may not be read by unauthorized devices.

[1568] In example embodiments, the distributed ledger module 29030 generates a block in response to a triggering event. Examples of triggering events may include a predetermined time (e.g., every minute, every hour, every day), when a potential issue is classified or predicted, when one or more sensor measurements are outside of a tolerance threshold, or the like. In response to the triggering event, the distributed ledger module 29030 may generate a block based on sensor data that is to be reported. Depending on the configuration of the server kit 28700 and the intended use of the distributed ledger 28762, the amount of data and type of data that is included in a block may vary. For example, in a manufacturing or resource extraction setting such as the manufacturing facility 1700 or the underwater industrial setting 1800, the distributed ledger 28762 may be used to demonstrate functional machinery and/or to predict maintenance needs. In this example, the distributed ledger module 29030 may be accessible by insurance providers to set insurance rates and/or issue refunds. Thus, in this example, the distributed ledger module 29030 may include any sensor measurements (and related metadata) that are outside of a tolerance threshold or instance where an issue is classified or predicted. In another example, the distributed ledger may be accessible by a regulatory body to ensure that a facility is operating in accordance with one or more regulations. In these embodiments, the distributed ledger module 29030 may store a set of one or

more sensor measurements (and related metadata) in a block, such that the sensor measurements may be analyzed by the regulatory agency. In some of these embodiments, the sensor measurements may be compressed to store more sensor data in a single block. In response to generating a block, the distributed ledger module 29030 may transmit the block to one or more node computing devices 28760. Upon the block being verified (e.g., using a consensus mechanism), each node computing device 28760 may update the distributed ledger 28762 with the new block.

[1569] As discussed in some embodiments, the distributed ledger may further include smart contracts. Once written, a smart contract may be encoded in a block and deployed to the distributed ledger 28762. The address of the smart contract (e.g., the block ID of the block containing the smart contract) may be provided to one or more parties to the smart contract, such that respective parties may invoke the smart contract using the address. In some of these embodiments, the address of the smart contract may be provided to the distributed ledger module 29030, such that the distributed ledger module 29030 may report items to the smart contract. In example embodiments, the distributed ledger module 29030 may leverage the API of a smart contract to report the items to the smart contract.

[1570] In example implementations discussed above, an insurer may utilize a smart contract to allow insured facility owners and/or operators to demonstrate that the equipment in the facility is functioning properly. In example embodiments, the smart contract may trigger the issuance of rebates or refunds on portions of insurance premiums when an owner and/or operator of a facility provides sufficient sensor data that indicates the facility is operating without issue. In some of these embodiments, the smart contract may include a first condition that requires a certain amount of sensor data to be reported by a facility and a second condition that each instance of the sensor data equals a value (e.g., no classified or predicted issues) or range of values (e.g., all sensor measurements within a predefined range of values). In example embodiments, the action may be to deposit funds (e.g., a wire transfer or cryptocurrency) into an account in response to the first and second conditions being met. In this example, the distributed ledger module 29030 may write blocks containing sensor data to the distributed ledger 28762. The distributed ledger module 29030 may also provide the addresses of these blocks to the smart contract (e.g., via an API of the smart contract). Upon the smart contract verifying the first and second conditions of the contract, the smart contract may initiate the transfer of funds from an account of the insurer to the account of the insured.

[1571] In another example discussed above, a regulatory body (e.g., a state, local, or federal regulatory agency) may utilize a smart contract that monitors facilities (e.g., food inspection facilities, pharmaceutical manufacturing facilities, indoor agricultural facilities, offshore oil extraction facilities, or the like) based on reported sensor data to ensure compliance with one or more regulations. In example embodiments, the smart contract may be configured to receive and verify the sensor data from a facility (e.g., via an API of the smart contract), and in response to verifying the sensor data issues a compliance token (or certificate) to an account of the facility owner. In some of these embodiments, the smart contract may include a first condition that requires

a certain amount of sensor data to be reported by a facility and a second condition that requires the sensor data to be compliant with the reporting regulations. In this example, the distributed ledger module 29030 may write blocks containing sensor data to the distributed ledger. The sensor kit 28700 may also provide the addresses of these blocks to the smart contract (e.g., using an API of the smart contract). Upon the smart contract verifying the first and second conditions of the contract, the smart contract may generate a token indicating compliance by the facility operator, and may initiate the transfer of funds to an account (e.g., a digital wallet) associated with the facility.

[1572] FIG. 291 illustrates an example backend system 28750 according to some embodiments of the present disclosure. In example embodiments, the backend system 28750 may be implemented as a cloud service that is executed at one or more physical server devices. In example embodiments, the backend system 28750 may include a storage system 29102, a communication system 29104, and a processing system 29106. The backend system 28750 may include additional components not shown.

[1573] A storage system 29102 includes one or more storage devices. The storage devices may include persistent storage mediums (e.g., flash memory drive, hard disk drive) and/or transient storage devices (e.g., RAM). The storage system 29102 may store one or more data stores. A data store may include one or more databases, tables, indexes, records, filesystems, folders and/or files. In the illustrated embodiments, the storage system 29102 stores a sensor kit data store 29110 and a model data store 29112. A storage system 29102 may store additional or alternative data stores without departing from the scope of the disclosure.

[1574] In example embodiments, the sensor kit data store 29110 stores data relating to respective sensor kits 28700. In example embodiments, the sensor kit data store 29110 may store sensor kit data corresponding to each installed sensor kit 28700. In example embodiments, the sensor kit data may indicate the devices in a sensor kit 28700, including each sensor 28702 (e.g., a sensor ID) in the sensor kit 28700. In example embodiments, the sensor kit data may indicate the sensor data captured by the sensor kit 28700. In some of these embodiments, the sensor kit data may identify each instance of sensor data captured by the sensor kit 28700, and for each instance of sensor data, the sensor kit data may indicate the sensor 28702 that captured the sensor data and, in example embodiments, a time stamp corresponding to the sensor data.

[1575] In example embodiments, the model data store 29112 stores machine-learned models that are trained by the AI system 29124 based on training data. The machine-learned models may include prediction models and classification models. In example embodiments, the training data used to train a particular model includes data collected from one or more sensor kits 28700 that monitor the same type of industrial setting 28720. The training data may additionally or alternatively may include historical data and/or expert generated data. In example embodiments, each machine-learned model may pertain to a respective type of industrial setting 28720. In some of these embodiments, the AI system 29124 may periodically update a machine-learned model pertaining to a type of industrial setting 28720 based on sensor data collected from sensor kits 28700 monitoring those types of industrial setting 28720 and outcomes obtained from those

industrial setting 28720. In example embodiments, machine-learned models pertaining to a type of industrial setting 28720 may be provided to the edge devices 28704 of sensor kits 28700 monitoring that type of industrial setting 28720.

[1576] In example embodiments, a communication system 29104 includes one or more communication devices, including at least one external communication device that communicates with a public communication network (e.g., the Internet) ether. The external communication devices may perform wired or wireless communication. In example embodiments, the external communication devices may include cellular chipsets (e.g., 4G or 5G chipsets), Ethernet cards and/or Wi-Fi cards, or other suitable communication devices.

[1577] In example embodiments, the processing system 29106 may include one or more memory devices (e.g., ROM and/or RAM) that store computer-executable instructions and one or more processors that execute the computer-executable instructions. The processors may execute in a parallel or distributed manner. The processors may be located in the same physical server device or in different server devices. The processing system 29106 may execute one or more of a decoding module 29120, a data processing module 29122, an AI module 29124, a notification module 29126, an analytics module 29128, a control module 29130, a dashboard module 29132, a configuration module 29134, and a distributed ledger management module 29136. The processing system 406 may execute additional or alternative modules without departing from the scope of the disclosure. Furthermore, the modules discussed herein may include submodules that perform one or more functions of a respective module.

[1578] In example embodiments, a sensor kit 28700 may transmit encoded sensor kit packets containing sensor data to the backend system 28750. In these embodiments, the decoding module 29120 may receive encoded sensor data from an edge device 28704 and may decrypt, decode, and/or decompress the encoded sensor kit packets to obtain the sensor data and metadata relating to the received sensor data (e.g., a sensor kit id and one or more sensor ids of sensors that captured the sensor data). The decoding module 29120 may output the sensor data and any other metadata to the data processing module 29122.

[1579] In example embodiments, the data processing module 29122 may process the sensor data received from the sensor kits 28700. In example embodiments, the data processing module 29122 may receive the sensor data and may store the sensor data in the sensor kit data store 29110 in relation to the sensor kit 28700 that provided to the sensor data. In example embodiments, the data processing system 29122 may provide AI-related requests to the AI module 29124. In these embodiments, the data processing system 29122 may extract relevant sensor data instances from the received sensor data and may provide the extracted sensor data instances to the AI module 29124 in a request that indicates the type of request (e.g., what type of prediction or classification) and the sensor data to be used. In the event a potential issue is predicted or classified, the data processing module 29122 may execute a workflow associated with the potential issue. A workflow may define the manner by which a potential issue is handled. For instance, the workflow may indicate that a notification should be transmitted to a human user, a remedial action should be

initiated, and/or other suitable actions. The data processing module 29122 may perform additional or alternative processing tasks without departing from the scope of the disclosure.

[1580] In example embodiments, the AI module 29124 trains machine-learned models that are used to make predictions or classifications. The machine-learned models may include any suitable type of models, including neural networks, deep neural networks, recursive neural networks, Bayesian neural networks, regression-based models, decision trees, prediction trees, classification trees, Hidden Markov Models, and/or any other suitable types of models. The AI module 29124 may train a machine-learned model on a training data set. A training data set may include expert-generated data, historical data, and/or outcome-based data. Outcome-based data may be data that is collected after a prediction or classification is made that indicates whether the prediction or classification was correct or incorrect and/or a realized outcome. A training data instance may refer to a unit of training data that includes a set of features and a label. In example embodiments, the label in a training data instance may indicate a condition of an industrial component or an industrial setting 28720 at a given time. Examples of conditions will vary greatly depending on the industrial setting 28720 and the conditions that the machine-learning model is being trained to predict or classify. Examples of labels in a manufacturing facility may include, but are not limited to, no issues detected, a mechanical failure of a component, an electrical failure of a component, a chemical leak detected, and the like. Examples of labels in a mining facility may include, but are not limited to, no issues detected, an oxygen deficiency, the presence of a toxic gas, a failing structural component, and the like. Examples of labels in an oil and/or gas facility (e.g., oil field, gas field, oil refinery, pipeline) may include, but are not limited to, no issues detected, a mechanical failure of a component (e.g., a failed valve or failed O-ring), a leak, and the like. Examples of labels in an indoor agricultural facility may include, but are not limited to, no issues detected, a plant died, a plant wilted, a plant turned a certain color (e.g., brown, purple, orange, or yellow), mold found, and the like. In each of these examples, there are certain features that may be relevant to a condition and some features that may have little or no bearing on the condition. In example embodiments, the AI module 29124 may reinforce the machine-learned models as more sensor data and outcomes relating to the machine-learned models are received. In example embodiments, the machine-learned models may be stored in the model data store 29112. Each model may be stored with a model identifier, which may be indicative of (e.g., mapped to) the type of industrial setting 28720 that the model makes, the type of prediction or classification made by the model, and the features that the model receives. In example embodiments, one or more machine-learned models (and subsequent updates thereto) may be pushed to respective sensor kits 28700, whereby the edge devices 28704 of the respective sensor kits 28700 may use one or more machine-learned model to make predictions and/or classifications without having to rely on the backend system 28750.

[1581] In example embodiments, the AI module 29124 receives requests for predictions and/or classifications and determines predictions and/or classifications based on the requests. In example embodiments, a request may indicate a type of prediction or classification that is being requested and may include a set of features for making the prediction or classification. In response to the

request, the AI module 29124 may select a machine-learned model to leverage based on the type of prediction or classification being requested, whereby the selected model receives a certain set of features. The AI module 29124 may then generate a feature vector that includes one or more instances of sensor data and may feed the feature vector into the selected model. In response to the feature vector, the selected model may output a prediction or classification, and a degree of confidence (e.g., a confidence score) in the prediction or classification. The AI module 29124 may output the prediction or classification, as well as the degree of confidence therein, to the module that provided the request.

[1582] In example embodiments, the notification module 29126 may issue notifications to users and/or respective industrial setting 28720 when an issue is detected in a respective setting. In example embodiments, a notification may be sent to a user device of a user indicating the nature of the issue. The notification module 29126 may implement an API (e.g., a REST API), whereby a user device of a user associated with the industrial setting 28720 may request notifications from the backend system 28750. In response to the request, the notification module 29126 may provide any notifications, if any, to the user device. In example embodiments, a notification may be sent to a device located at an industrial setting 28720, whereby the device may raise an alarm at the industrial setting 28720 in response to the industrial setting 28720.

[1583] In example embodiments, the analytics module 29128 may perform analytics related tasks on sensor data collected by the backend system 28750 and stored in the sensor kit data store 29110. In example embodiments, the analytics tasks may be performed on sensor data received from individual sensor kits. Additionally, or alternatively, the analytics tasks may be performed on sensor data Examples of analytics tasks that may be performed on sensor data obtained from various sensor kits 28700 monitoring different industrial setting 28720. Examples of analytics tasks may include energy utilization analytics, quality analytics, process optimization analytics, financial analytics, predictive analytics, yield optimization analytics, fault prediction analytics, scenario planning analytics, and many others.

[1584] In example embodiments, the control module 29130 may control one or more aspects of an industrial setting 28720 based on a determination made by the AI system 29124. In example embodiments, the control module 29130 may be configured to provide commands to a device or system at the industrial setting 28720 to take a remedial action in response to a particular issue being detected. For example, the control module 29130 may issue a command to a manufacturing facility to stop an assembly line in response to a determination that a critical component on the assembly line is likely failing or likely failed. In another example, the control module 29130 may issue a command to an agricultural facility to activate a dehumidifier in response to a determination that the humidity levels are too high in the facility. In another example, the control module 29130 may issue a command to shut a valve in an oil pipeline in response to a determination that a component in the oil pipeline downstream to the valve is likely failing or likely failed. For a particular industrial setting 28720, the control module 29130 may perform remedial actions defined by a human user associated with the industrial setting 28720, such that the human user may define what conditions may trigger the remedial action.

[1585] In example embodiments, the dashboard module 29132 presents a dashboard to human users via a user device 28740 associated with the human user. In example embodiments, the dashboard provides a graphical user interface that allows the human user to view relating to a sensor kit 28700 with which the human user is associated (e.g., an employee at the industrial setting 28720). In these embodiments, the dashboard module 29132 may retrieve and display raw sensor data provided by the sensor kit, analytical data relating to the sensor data provided by the sensor kit 28700, predictions or classifications made by the backend system 28750 based on the sensor data, and the like.

[1586] In example embodiments, the dashboard module 29132 allows human users to configure aspects of the sensor kits 28700. In example embodiments, the dashboard module 29132 may present a graphical user interface that allows a human user to configure one or more aspects of a sensor kit 28700 with which the human user is associated. In example embodiments, the dashboard may allow a user to configure alarm limits with respect to one or more sensor types and/or conditions. For example, a user may define a temperature value at which a notification is sent to a human user. In another example, the user may define a set of conditions, which if predicted by the AI module and/or the edge device, trigger an alarm. In example embodiments, the dashboard may allow a user to define which users receive a notification when an alarm is triggered. In example embodiments, the dashboard may allow a user to subscribe to additional features of the backend system 28750 and/or an edge device 28704.

[1587] In example embodiments, the dashboard may allow a user to add one or more subscriptions to a sensor kit 28700. The subscriptions may include access to backend services and/or edge services. A user may select a service to add to a sensor kit 28700 and may provide payment information to pay for the services. Upon verification of the payment information, the backend system 28750 may provide the sensor kit 28700 access to those features. Examples of services that may be subscribed to include analytics services, AI-services, notification services, and the like. The dashboard may allow the user to perform additional or alternative configurations.

[1588] In example embodiments, the configuration module 29134 maintains configurations of respective sensor kits 28700. Initially, when a new sensor kit 28700 is deployed in an industrial setting 28720, the configuration module 29134 may update the sensor kit data store 29110 with the device IDs of each device in the newly installed sensor kit 28700. Once the sensor kit data store 29110 has updated the sensor kit data store 29110 to reflect the newly installed sensor kit 28700, the backend system 28750 may begin storing sensor data from the sensor kit 28700. In example embodiments, new sensors 28702 may be added to respective sensor kits 28700. In these embodiments, an edge device 28704 may provide an add request to the backend system 28750 upon an attempt to add a device to the sensor kit 28700. In example embodiments, the request may indicate a sensor ID of the new sensor. In response to the request, the configuration module 29134 may add the sensor ID of the new sensor to the sensor kit data of the requesting sensor kit 28700 in the sensor kit data store 29110.

[1589] In example embodiments, the backend system 28750 includes a distributed ledger management module 29136. In some of these embodiments, the distributed ledger management

module 29136 allows a user to update and/or configure a distributed ledger. In some of these embodiments, the distributed ledger management module 29136 allows a user to define or upload a smart contract. As discussed, the smart contract may include one or more conditions that are verified by the smart contract and one or more actions that are triggered when the conditions are verified. In example embodiments, the user may provide one or more conditions that are to be verified to the distributed ledger management module 29136 via a user interface. In some of these embodiments, the user may provide the code (e.g., JavaScript code, Java code, C code, C++ code, etc.) that defines the conditions. The user may also provide the actions that are to be performed in response to certain conditions being met. In response to a smart contract being uploaded/created, the distributed ledger management module 29136 may deploy the smart contract. In example embodiments, the distributed ledger management module 29136 may generate a block containing the smart contract. The block may include a header that defines an address of the block, and a body that includes an address to a previous block and the smart contract. In example embodiments, the distributed ledger management module 29136 may determine a hash value based on the body of the block and/or may encrypt the block. The distributed ledger management module 29136 may transmit the block to one or more node computing devices 28760, which in turn update the distributed ledger with the block containing the smart contract. The distributed ledger management module 29136 may further provide the address of the block to one or more parties that may access the smart contract. The distributed ledger management module 29136 may perform additional or alternative functions without departing from the scope of the disclosure.

[1590] The backend system 28750 may include additional or alternative components, data stores, and/or modules that are not discussed.

[1591] FIG. 292 illustrates an example set of operations of a method 29200 for compressing sensor data obtained by a sensor kit 28700. In example embodiments, the method 29200 may be performed by an edge device 28704 of a sensor kit 28700.

[1592] At 29210, the edge device 28704 receives sensor data from one or more sensors 28702 of the sensor kit 28700 via a sensor kit network 200. In example embodiments, the sensor data from a respective sensor 28702 may be received in a reporting packet. Each reporting packet may include a device identifier of the sensor 28702 that generated the reporting packet and one or more instances of sensor data captured by sensor 28702. The reporting packet may include additional data, such as a timestamp or other metadata.

[1593] At 29212, the edge device 28704 processes the sensor data. In example embodiments, the edge device 28704 may dedupe any reporting packets that are duplicative. In example embodiments, the edge device 28704 may filter out sensor data that is clearly erroneous (e.g., outside of a tolerance range). In example embodiments, the edge device 28704 may aggregate the sensor data obtained from multiple sensors 28702. In example embodiments, the edge device 28704 may perform one or more AI related tasks, such as determining a prediction or classification relating to a condition of one or more industrial components of the industrial setting 28720. In some of these embodiments, the decision to compress the sensor data may depend on whether the edge device 28704 determines that there are any potential issues with the industrial component.

For example, the edge device 28704 may compress the sensor data when there have been no issues predicted or classified. In other embodiments, the edge device 28704 may compress any sensor data that is being transmitted to the backend system or certain types of sensor data (e.g., sensor data obtained from temperature sensors).

[1594] At 29214, the edge device 28704 may compress the sensor data. The edge device 28704 may employ any suitable compression techniques for compressing the sensor data. For example, the edge device 28704 may employ vertical or horizontal compression techniques. The edge device 28704 may be configured with a codec that compresses the sensor data. The codec may be a proprietary codec or an "off-the-shelf" codec.

[1595] At 29216, the edge device 28704 may transmit the compressed sensor data to the backend system 28750. In example embodiments, the edge device 28704 may generate a sensor kit packet that contains the compressed data. The sensor kit packet may designate the source of the sensor kit packet (e.g., a sensor kit ID or edge device ID) and may include additional metadata (e.g., a timestamp). In example embodiments, the edge device 28704 may encrypt the sensor kit packet prior to transmitting the sensor kit packet to the backend system 28750. In example embodiments, the edge device 28704 transmits the sensor kit packet to the backend system 28750 directly (e.g., via a cellular connection, a network connection, or a satellite uplink). In other embodiments, the edge device 28704 transmits the sensor kit packet to the backend system 28750 via a gateway device, which transmits the sensor kit packet to the backend system 28750 directly (e.g., via a cellular connection or a satellite uplink).

[1596] FIG. 293 illustrates an example set of operations of a method 29300 for processing compressed sensor data received from a sensor kit 28700. In example embodiments, the method 29300 is executed by a backend system 28750.

[1597] At 29310, the backend system 28750 receives compressed sensor data from a sensor kit. In example embodiments, the compressed sensor data may be received in a sensor kit packet.

[1598] At 29312, the backend system 28750 decompresses the received sensor data. In example embodiments, the backend system may utilize a codec to decompress the received sensor data. Prior to decompressing the received sensor data, the backend system 28750 may decrypt a sensor kit packet containing the compressed sensor data.

[1599] At 29314, the backend system 28750 performs one or more backend operations on the decompressed sensor data. The backend operations may include storing the data, filtering the data, performing AI-related tasks on the sensor data, issuing one or more notifications in relation to the results of the AI-related tasks, performing one or more analytics related tasks, controlling an industrial component of the industrial setting 28720, and the like.

[1600] FIG. 294 illustrates an example set of operations of a method 29400 for streaming sensor data from a sensor kit 28700 to a backend system 28750. In example embodiments, the method 29400 may be executed by an edge device 28704 of the sensor kit 28700.

[1601] At 29410, the edge device 28704 receives sensor data from one or more sensors 28702 of the sensor kit 28700 via a sensor kit network 28800. In example embodiments, the sensor data from a respective sensor 28702 may be received in a reporting packet. Each reporting packet may

include a device identifier of the sensor 28702 that generated the reporting packet and one or more instances of sensor data captured by sensor 28702. The reporting packet may include additional data, such as a timestamp or other metadata. In example embodiments, the edge device 28704 may process the sensor data. For example, the edge device 28704 may dedupe any reporting packets that are duplicative and/or may filter out sensor data that is clearly erroneous (e.g., outside of a tolerance range). In example embodiments, the edge device 28704 may aggregate the sensor data obtained from multiple sensors 28702.

[1602] At 29412, the edge device 28704 may normalize and/or transform the sensor data into a media-frame compliant format. In example embodiments, the edge device 28704 may normalize and/or transform each sensor data instance into a value that adheres to the restrictions of a media frame that will contain the sensor data. For example, in embodiments where the media frames are video frames, the edge device 28704 may normalize and/or transform instances of sensor data into acceptable pixel frames. The edge device 28704 may employ one or more mappings and/or normalization functions to transform and/or normalize the sensor data.

[1603] At 29414, the edge device 28704 may generate a block of media frames based on the transformed and/or normalized sensor data. For example, in embodiments where the media frames are video frames, the edge device 28704 may populate each instance of transformed and/or normalized sensor data into a respective pixel of the video frame. The manner by which the edge device 28704 assigns an instance of transformed and/or normalized sensor data to a respective pixel may be defined in a mapping that maps respective sensors to respective pixel values. In example embodiments, the mapping may be defined so as to minimize variance between the values in adjacent pixels. In example embodiments, the edge device 28704 may generate a series of time-sequenced media frames, such that each successive media frame corresponds to a subsequent set of sensor data instances.

[1604] At 29416, the edge device 28704 may encode the block of the media frame. In example embodiments, the edge device 28704 may employ an encoder of a media codec (e.g., a video codec) to compress the block of media frames. The codec may be a proprietary codec or an "off-the-shelf" codec. For example, the media codec may be an H.264/MPEG-4 codec, an H.265/MPEG-H codec, an H.263/MPEG-4 codec, proprietary codecs, and the like. The codec receives the block of media frames and generates an encoded media block based thereon.

[1605] At 29418, the edge device 28704 may transmit the encoded media block to the backend system 28750. In example embodiments, the edge device 28704 may stream the encoded media blocks to the backend system 28750. Each encoded block may designate the source of the block (e.g., a sensor kit ID or edge device ID) and may include additional metadata (e.g., a timestamp and/or a block identifier). In example embodiments, the edge device 28704 may encrypt the encoded media blocks prior to transmitting encoded media blocks to the backend system 28750. The edge device 28704 may transmit the encoded media blocks to the backend system 28750 directly (e.g., via a cellular connection, a network connection, or a satellite uplink) or via a gateway device, which transmits the encoded media block to the backend system 28750 directly (e.g., via a cellular connection or a satellite uplink).

[1606] The edge device 28704 may continue to execute the foregoing method 29400, so as to deliver a stream of live sensor data from a sensor kit. The foregoing method 29400 may be performed in settings where there are many sensors deployed within the setting and the sensors are sampled frequently or continuously. In this way, the bandwidth required to provide the sensor data to the backend system is reduced.

[1607] FIG. 295 illustrates an example set of operations of a method 29500 for ingesting a sensor data stream from an edge device 28704. In example embodiments, the method 29500 is executed by a backend system.

[1608] At 29510, the backend system 28750 receives an encoded media block from a sensor kit. The backend system 28750 may receive encoded media blocks as part of a sensor data stream.

[1609] At 29512, the backend system 28750 decodes the encoded block using a decoder corresponding to the codec of the codec used to encode the media block to obtain a set of successive media frames. As discussed with respect to the encoding operation, the codec may be a proprietary codec or an "off-the-shelf" codec. For example, the media codec may be an H.264/MPEG-4 codec, an H.265/MPEG-H codec, an H.263/MPEG-4 codec, proprietary codecs, and the like. The codec receives the encoded block of media frames and decodes the encoded block to obtain a set of sequential media frames.

[1610] At 29514, the backend system 28750 recreates the sensor data based on the media frame. In example embodiments, the backend system 28750 determines the normalized and/or transformed sensor values embedded in each respective media frame. For example, in embodiments where the media frames are video frames, the backend system 28750 may determine pixel values for each pixel in the media frame. A pixel value may correspond to respective sensor 28702 of a sensor kit 28700 and the value may represent a normalized and/transformed instance of sensor data. In example embodiments, the backend system 28750 may recreate the sensor data by inversing the normalization and/or transformation of the pixel value. In example embodiments, the backend system 28750 may utilize an inverse transformation and/or an inverse normalization function to obtain each recreated sensor data instance.

[1611] At 29518, the backend system 28750 performs one or more backend operations based on the recreated sensor data. The backend operations may include storing the data, filtering the data, performing AI-related tasks on the sensor data, issuing one or more notifications in relation to the results of the AI-related tasks, performing one or more analytics related tasks, controlling an industrial component of the industrial setting 28720, and the like.

[1612] FIG. 296 illustrates a set of operations of a method 29600 for determining a transmission strategy and/or a storage strategy for sensor data collected by a sensor kit 28700 based on the sensor data. A transmission strategy may define a manner that sensor data is transmitted (if at all) to the backend system. For example, sensor data may be compressed using an aggressive lossy codec, compressed using a lossless codec, and/or transmitted without compression. A storage strategy may define a manner by which sensor data is stored at the edge device 28704. For example, sensor data may be stored permanently (or until a human removes the sensor data), may be stored for a period of time (e.g., one year) or may be discarded. The method 29600 may be executed by an

edge device 28704. The method 29600 may be executed to reduce the network bandwidth consumed by the sensor kit 28700 and/or reduce the storage constraints at the edge device 28704.

[1613] At 29610, the edge device 28704 receives sensor data from the sensors 28702 of the sensor kit 28700. The data may be received continuously or intermittently. In example embodiments, the sensors 28702 may push the sensor data to the edge device 28704 and/or the edge device 28704 may request the sensor data 28702 from the sensors 28702 periodically. In example embodiments, the edge device 28704 may process the sensor data upon receipt, including deduping the sensor data.

[1614] In example embodiments, the edge device 28704 may be configured to perform one or more AI-related tasks prior to transmission via the satellite uplink. In some of these embodiments, the edge device 28704 may be configured to determine whether there are likely no issues relating to any of the components and/or the industrial setting 28720 based on the sensor data and one or more machine-learned models.

[1615] At 29612, the edge device 28704 may generate one or more feature vectors based on the sensor data. The feature vectors may include sensor data from a single sensor 28702, a subset of sensors 28702, or all of the sensors 28702 of the sensor kit 28700. In scenarios where a single sensor or a subset of sensors 28702 are included in the feature vector, the machine-learned model may be trained to identify one or more issues relating to an industrial component or the industrial setting 28720, but may not be sufficient to fully deem the entire setting as likely safe/free from issues. Additionally or alternatively, the feature vectors may correspond to a single snapshot in time (e.g., all sensor data in the feature vector corresponds to the same sampling event) or over a period of time (sensor data samples from a most recent sampling event and sensor data samples from previous sampling events). In embodiments where the feature vectors define sensor data from a single snapshot, the machine-learned models may be trained to identify potential issues without any temporal context. In embodiments where the feature vectors define sensor data over a period of time, the machine-learned models may be trained to identify potential issues with the context of what the sensor(s) 28702 was/were reporting previously. In these embodiments, the edge device 28704 may maintain a cache of sensor data that is sampled over a predetermined time (e.g., previous hour, previous day, previous N days), such that the cache is cleared out in a first-in-first-out manner. In these embodiments, the edge device 28704 may retrieve the previous sensor data samples from the cache to use to generate feature vectors that have data samples spanning a period of time.

[1616] At 29614, the edge device 28704 may input the one or more feature vectors into one or more respective machine-learned models. A respective model may output a prediction or classification relating to an industrial component and/or the industrial setting 28720, and a confidence score relating to the prediction or classification.

[1617] At 29616, the edge device 28704 may determine a transmission strategy and/or a storage strategy based on the output of the machine-learned models. In example embodiments, the edge device 28704 may make determinations relating to the manner by which sensor data is transmitted to the backend system 28750. In example embodiments, the edge device 28704 may make

determinations relating to the manner by which sensor data is transmitted to the backend system 28750 and/or stored at the edge device. In some of these embodiments, the edge device 28704 may compress sensor data when there are no likely issues across the entire industrial setting 28720 and individual components of the industrial setting 28720. For example, if the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence (e.g., the confidence score is greater than .98), the edge device 28704 may compress the sensor data. Alternatively, in the scenario where the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence, the edge device 28704 may forego transmission but may store the sensor data at the edge device 28704 for a predefined period of time (e.g., a one-year expiry). In scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may transmit the sensor data without compressing the sensor data or using a lossless compression codec. Additionally or alternatively, in scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may store the sensor data used to make the prediction or classification indefinitely, as well as data that was collected prior to and/or after the condition was predicted or classified.

[1618] FIG. 297 illustrates an example configuration of a sensor kit 29700 according to some embodiments of the present disclosure. In the illustrated example, the sensor kit 29700 is configured to communicate with a communication network 28780 via an uplink 29708 to a satellite 29710. In example embodiments, the sensor kit 29700 of FIG. 151 is configured for use in industrial setting 28720 located in remote locations, where cellular coverage is unreliable or non-existent. In example embodiments, the sensor kit 29700 may be installed in natural resource extraction, natural resource transportation systems, power generation facilities, and the like. For example, the sensor kit 29700 may be deployed in an oil or natural gas fields, off-shore oil rigs, mines, oil or gas pipelines, solar fields, wind farms, hydroelectric power stations, and the like.

[1619] In the example of FIG. 151, the sensor kit 29700 includes an edge device 28704 and a set of sensors 28702. The sensors 28702 may include various types of sensors 28702, which may vary depending on the industrial setting 28720. In the illustrated example, the sensors 28702 communicate with the edge device 28704 via a mesh network. In these embodiments, the sensors 28702 may communicate sensor data to proximate sensors 28702, so as to propagate the sensor data to the edge device 28704 located at the remote/peripheral areas of the industrial setting 28720 to the edge device 28704. While a mesh network is shown, the sensor kits 29700 of FIG. 151 may include alternative network topologies, such as a hierarchal topology (e.g., some or all of the sensors 28702 communicate with the edge device 28704 via respective collection devices) or a star topology (e.g., sensors 28702 communicate to the edge device directly).

[1620] In the embodiments of FIG. 151, the edge device 28704 includes a satellite terminal with a directional antenna that communicates with a satellite. The satellite terminal may be pre-configured to communicate with a geosynchronous or low Earth orbit satellites. The edge device 28704 may receive sensor data from the sensor kit network established by the sensor kit 29700. The edge

device 28704 may then transmit the sensor data to the backend system 28750 via the satellite 29710.

[1621] In example embodiments, the configurations of the sensor kit 29700 are suited for industrial setting 28720 covering a remote area where external power sources are not abundant. In example embodiments, the sensor kit 29700 may include external power sources, such as batteries, rechargeable batteries, generators, and/or solar panels. In these embodiments, the external power sources may be deployed to power the sensors 28702, the edge device 28704, and any other devices in the sensor kit 29700.

[1622] In example embodiments, the configurations of the sensor kit 29700 are suited for outdoor industrial setting 28720. In example embodiments, the sensors 28702, the edge device 28704, and other devices of the sensor kit 28700 (e.g., collection devices) may be configured with weatherproof housings. In these embodiments, the sensor kit 29700 may be deployed in an outdoor setting.

[1623] In example embodiments, the edge device 28704 may be configured to perform one or more AI-related tasks prior to transmission via the satellite uplink. In some of these embodiments, the edge device 28704 may be configured to determine whether there are likely no issues relating to any of the components and/or the industrial setting 28720 based on the sensor data and one or more machine-learned models. In example embodiments, the edge device 28704 may receive the sensor data from the various sensors and may generate one or more feature vectors based thereon. The feature vectors may include sensor data from a single sensor 28702, a subset of sensors 28702, or all of the sensors 28702 of the sensor kit 29700. In scenarios where a single sensor or a subset of sensors 28702 are included in the feature vector, the machine-learned model may be trained to identify one or more issues relating to an industrial component or the industrial setting 28720, but may not be sufficient to fully deem the entire setting as likely safe/free from issues. Additionally or alternatively, the feature vectors may correspond to a single snapshot in time (e.g., all sensor data in the feature vector corresponds to the same sampling event) or over a period of time (sensor data samples from a most recent sampling event and sensor data samples from previous sampling events). In embodiments where the feature vectors define sensor data from a single snapshot, the machine-learned models may be trained to identify potential issues without any temporal context. In embodiments where the feature vectors define sensor data over a period of time, the machine-learned models may be trained to identify potential issues with the context of what the sensor(s) 28702 was/were reporting previously. In these embodiments, the edge device 28704 may maintain a cache of sensor data that is sampled over a predetermined time (e.g., previous hour, previous day, previous N days), such that the cache is cleared out in a first-in-first-out manner. In these embodiments, the edge device 28704 may retrieve the previous sensor data samples from the cache to use to generate feature vectors that have data samples spanning a period of time.

[1624] In example embodiments, the edge device 28704 may feed the one or more feature vectors into one or more respective machine-learned models. A respective model may output a prediction or classification relating to an industrial component and/or the industrial setting 28720, and a confidence score relating to the prediction or classification. In example embodiments, the edge

device 28704 may make determinations relating to the manner by which sensor data is transmitted to the backend system 28750 and/or stored at the edge device. For instance, in example embodiments, the edge device 28704 may compress sensor data based on the prediction or classification. In some of these embodiments, the edge device 28704 may compress sensor data when there are no likely issues across the entire industrial setting 28720 and individual components of the industrial setting 28720. For example, if the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence (e.g., the confidence score is greater than .98), the edge device 28704 may compress the sensor data. Alternatively, in the scenario where the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence, the edge device 28704 may forego transmission but may store the sensor data at the edge device 28704 for a predefined period of time (e.g., one year). In scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may transmit the sensor data without compressing the sensor data or using a lossless compression codec. In this way, the amount of bandwidth that is transmitted via the satellite uplink may be reduced, as the majority of the time the sensor data will be compressed or not transmitted.

[1625] In example embodiments, the edge device 28704 may apply one or more rules to determine whether a triggering condition exists. In example embodiments, the one or more rules may be tailored to identify potentially dangerous and/or emergency situations. In these embodiments, the edge device 28704 may trigger one or more notifications or alarms when a triggering condition exists. Additionally or alternatively, the edge device 28704 may transmit the sensor data without any compression when a triggering condition exists.

[1626] FIG. 298 illustrates an example configuration of a sensor kit 29800 according to some embodiments of the present disclosure. In the illustrated example, the sensor kit 29800 is configured to include a gateway device 29806 that communicates with a communication network 28780 via an uplink 29708 to a satellite 29710. In example embodiments, the sensor kit 29800 of FIG. 152 is configured for use in industrial setting 28720 located in remote locations, where cellular coverage is unreliable or non-existent, and where the edge device 28704 is located in a location where physical transmission to a satellite is unreliable or impossible. In example embodiments, the sensor kit 29700 may be installed in underground or underwater facilities, or in facilities having very thick walls. For example, the sensor kit 29700 may be deployed in underground mines, underwater oil or gas pipelines, underwater hydroelectric power stations, and the like.

[1627] In the example of FIG. 152, the sensor kit 29800 includes an edge device 28704, a set of sensors 28702, and a gateway device 29806. In example embodiments, the gateway device 29806 is a communication device that includes a satellite terminal with a directional antenna that communicates with a satellite. The satellite terminal may be pre-configured to communicate with a geosynchronous or low Earth orbit satellites. In example embodiments, the gateway device 29806 may communicate with the edge device 28704 via a wired communication link 29808 (e.g., Ethernet). The edge device 28704 may receive sensor data from the sensor kit network established

by the sensor kit 29800. The edge device 28704 may then transmit the sensor data to the gateway device 29806 via the wired communication link 29808. The gateway device 29806 may then communicate the sensor data to the backend system 28750 via the satellite uplink 29708.

[1628] The sensors 28702 may include various types of sensors 28702, which may vary depending on the industrial setting 28720. In the illustrated example, the sensors 28702 communicate with the edge device 28704 via a mesh network. In these embodiments, the sensors 28702 may communicate sensor data to proximate sensors 28702, so as to propagate the sensor data to the edge device 28704 located at the remote/peripheral areas of the industrial setting 28720 to the edge device 28704. While a mesh network is shown, the sensor kits 29800 of FIG. 152 may include alternative network topologies, such as a hierarchal topology (e.g., some or all of the sensors 28702 communicate with the edge device 28704 via respective collection devices) or a star topology (e.g., sensors 28702 communicate to the edge device directly).

[1629] In example embodiments, the configurations of the server kit 29800 are suited for industrial setting 28720 covering a remote area where external power sources are not abundant. In example embodiments, the sensor kit 29800 may include external power sources, such as batteries, rechargeable batteries, generators, and/or solar panels. In these embodiments, the external power sources may be deployed to power the sensors 28702, the edge device 28704, and any other devices in the sensor kit 29800.

[1630] In example embodiments, the configurations of the server kit 29800 are suited for underground or underwater industrial setting 28720. In example embodiments, the sensors 28702, the edge device 28704, and other devices of the sensor kit 28700 (e.g., collection devices) may be configured with waterproof housings or otherwise airtight housings (to prevent dust from entering the edge device 28704 and/or sensor devices 28702). Furthermore, as the gateway device 29808 is likely to be situated outdoors, the gateway device 29808 may include a weatherproof housing.

[1631] In example embodiments, the edge device 28704 may be configured to perform one or more AI-related tasks prior to transmission via the satellite uplink. In some of these embodiments, the edge device 28704 may be configured to determine whether there are likely no issues relating to any of the components and/or the industrial setting 28720 based on the sensor data and one or more machine-learned models. In example embodiments, the edge device 28704 may receive the sensor data from the various sensors and may generate one or more feature vectors based thereon. The feature vectors may include sensor data from a single sensor 28702, a subset of sensors 28702, or all of the sensors 28702 of the sensor kit 29800. In scenarios where a single sensor or a subset of sensors 28702 are included in the feature vector, the machine-learned model may be trained to identify one or more issues relating to an industrial component or the industrial setting 28720, but may not be sufficient to fully deem the entire setting as likely safe/free from issues. Additionally or alternatively, the feature vectors may correspond to a single snapshot in time (e.g., all sensor data in the feature vector corresponds to the same sampling event) or over a period of time (sensor data samples from a most recent sampling event and sensor data samples from previous sampling events). In embodiments where the feature vectors define sensor data from a single snapshot, the machine-learned models may be trained to identify potential issues without any temporal context.

In embodiments where the feature vectors define sensor data over a period of time, the machine-learned models may be trained to identify potential issues with the context of what the sensor(s) 28702 was/were reporting previously. In these embodiments, the edge device 28704 may maintain a cache of sensor data that is sampled over a predetermined time (e.g., previous hour, previous day, previous N days), such that the cache is cleared out in a first-in-first-out manner. In these embodiments, the edge device 28704 may retrieve the previous sensor data samples from the cache to use to generate feature vectors that have data samples spanning a period of time.

[1632] In example embodiments, the edge device 28704 may feed the one or more feature vectors into one or more respective machine-learned models. A respective model may output a prediction or classification relating to an industrial component and/or the industrial setting 28720, and a confidence score relating to the prediction or classification. In example embodiments, the edge device 28704 may make determinations relating to the manner by which sensor data is transmitted to the backend system 28750 and/or stored at the edge device. For instance, in example embodiments, the edge device 28704 may compress sensor data based on the prediction or classification. In some of these embodiments, the edge device 28704 may compress sensor data when there are no likely issues across the entire industrial setting 28720 and individual components of the industrial setting 28720. For example, if the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence (e.g., a confidence score is greater than .98), the edge device 28704 may compress the sensor data. Alternatively, in the scenario where the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence, the edge device 28704 may forego transmission but may store the sensor data at the edge device 28704 for a predefined period of time (e.g., one year). In scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may transmit the sensor data without compressing the sensor data or using a lossless compression codec. In this way, the amount of bandwidth that is transmitted via the satellite uplink may be reduced, as the majority of the time the sensor data will be compressed or not transmitted.

[1633] In example embodiments, the edge device 28704 may apply one or more rules to determine whether a triggering condition exists. In example embodiments, the one or more rules may be tailored to identify potentially dangerous and/or emergency situations. In these embodiments, the edge device 28704 may trigger one or more notifications or alarms when a triggering condition exists. Additionally or alternatively, the edge device 28704 may transmit the sensor data (via the gateway device 29806) without any compression when a triggering condition exists.

[1634] Figure 153 illustrates an example configuration of a sensor kit 29900 according to some embodiments of the present disclosure. In the example of figure 153, the sensor kit 29900 includes an edge device 28704, a set of sensors, and a set of collection devices. In example embodiments, the configurations of the sensor kit 29900 are suited for industrial setting 28720 covering a large area and where power sources are abundant; but where the industrial operator does not wish to connect the sensor kit 29900 to the private network of the industrial setting 28720. In example embodiments, the edge device 28704 includes a cellular communication device (e.g., a 4G LTE

chipset or 5G LTE chipset) with a transceiver that communicates with a cellular tower 29910. The cellular communication may be pre-configured to communicate with a cellular data provider. For example, in example embodiments, the edge device 28704 may include a SIM card that is registered with a cellular provider having a cellular tower 29910 that is proximate to the industrial setting 28720. The edge device 28704 may receive sensor data from the sensor kit network established by the sensor kit 29900. The edge device 28704 may process the sensor data and then transmit the sensor data to the backend system 28750 via the cellular tower 29910.

[1635] The sensors 28702 may include various types of sensors 28702, which may vary depending on the industrial setting 28720. In the illustrated example, the sensors 28702 communicate with the edge device 28704 via a hierarchical network. In these embodiments, the sensors 28702 may communicate sensor data to collection devices 206, which, in turn, may communicate the sensor data to edge device 28704 via a wired or wireless communication link. The hierarchical network may be deployed where the area being monitored is rather larger (e.g., over 40,000 sq. ft.) and power supplies are abundant, such as in a factory, a power plant, a food inspection facility, an indoor grow facility, and the like. While a hierarchal network is shown, the sensor kits 29900 of Figure 153 may include alternative network topologies, such as a mesh topology or a star topology (e.g., sensors 28702 communicate to the edge device directly).

[1636] In example embodiments, the edge device 28704 may be configured to perform one or more AI-related tasks prior to transmission via the satellite uplink. In some of these embodiments, the edge device 28704 may be configured to determine whether there are likely no issues relating to any of the components and/or the industrial setting 28720 based on the sensor data and one or more machine-learned models. In example embodiments, the edge device 28704 may receive the sensor data from the various sensors and may generate one or more feature vectors based thereon. The feature vectors may include sensor data from a single sensor 28702, a subset of sensors 28702, or all of the sensors 28702 of the sensor kit 29900. In scenarios where a single sensor or a subset of sensors 28702 are included in the feature vector, the machine-learned model may be trained to identify one or more issues relating to an industrial component or the industrial setting 28720, but may not be sufficient to fully deem the entire setting as likely safe/free from issues. Additionally or alternatively, the feature vectors may correspond to a single snapshot in time (e.g., all sensor data in the feature vector corresponds to the same sampling event) or over a period of time (sensor data samples from a most recent sampling event and sensor data samples from previous sampling events). In embodiments where the feature vectors define sensor data from a single snapshot, the machine-learned models may be trained to identify potential issues without any temporal context. In embodiments where the feature vectors define sensor data over a period of time, the machine-learned models may be trained to identify potential issues with the context of what the sensor(s) 28702 was/were reporting previously. In these embodiments, the edge device 28704 may maintain a cache of sensor data that is sampled over a predetermined time (e.g., previous hour, previous day, previous N days), such that the cache is cleared out in a first-in-first-out manner. In these embodiments, the edge device 28704 may retrieve the previous sensor data samples from the cache to use to generate feature vectors that have data samples spanning a period of time.

[1637] In example embodiments, the edge device 28704 may feed the one or more feature vectors into one or more respective machine-learned models. A respective model may output a prediction or classification relating to an industrial component and/or the industrial setting 28720, and a confidence score relating to the prediction or classification. In example embodiments, the edge device 28704 may make determinations relating to the manner by which sensor data is transmitted to the backend system 28750 and/or stored at the edge device. For instance, in example embodiments, the edge device 28704 may compress sensor data based on the prediction or classification. In some of these embodiments, the edge device 28704 may compress sensor data when there are no likely issues across the entire industrial setting 28720 and individual components of the industrial setting 28720. For example, if the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence (e.g., a confidence score is greater than .98), the edge device 28704 may compress the sensor data. Alternatively, in the scenario where the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence, the edge device 28704 may forego transmission but may store the sensor data at the edge device 28704 for a predefined period of time (e.g., one year). In scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may transmit the sensor data without compressing the sensor data or using a lossless compression codec. In this way, the amount of bandwidth that is transmitted via the cellular tower may be reduced, as the majority of the time the sensor data will be compressed or not transmitted.

[1638] In example embodiments, the edge device 28704 may apply one or more rules to determine whether a triggering condition exists. In example embodiments, the one or more rules may be tailored to identify potentially dangerous and/or emergency situations. In these embodiments, the edge device 28704 may trigger one or more notifications or alarms when a triggering condition exists. Additionally or alternatively, the edge device 28704 may transmit the sensor data without any compression when a triggering condition exists.

[1639] FIG. 154 illustrates an example configuration of a sensor kit 30000 according to some embodiments of the present disclosure. In the example of FIG. 154, the sensor kit 30000 includes an edge device 28704, a set of sensors 28702, a set of collection devices 206, and a gateway device 30006. In example embodiments, the configurations of the sensor kit 30000 are suited for industrial setting 28720 covering a large area and where power sources are abundant; but where the industrial operator does not wish to connect the sensor kit 30000 to the private network of the industrial setting 28720 and the walls of the industrial setting 28720 make wireless communication (e.g., cellular communication) unreliable or impossible. In example embodiments, the gateway device 30006 is a cellular network gateway device that includes a cellular communication device (e.g., 4G, 5G chipset) with a transceiver that communicates with a cellular tower 29910. The cellular communication may be pre-configured to communicate with a cellular data provider. For example, in example embodiments, the gateway device may include a SIM card that is registered with a cellular provider having a tower 29910 that is proximate to the industrial setting 28720. In example embodiments, the gateway device 30006 may communicate with the edge device 28704 via a wired

communication link 30008 (e.g., Ethernet). The edge device 28704 may receive sensor data from the sensor kit network established by the sensor kit 30000. The edge device 28704 may then transmit the sensor data to the gateway device 30006 via the wired communication link 30008. The gateway device 30006 may then communicate the sensor data to the backend system 28750 via the cellular tower 29910.

[1640] The sensors 28702 may include various types of sensors 28702, which may vary depending on the industrial setting 28720. In the illustrated example, the sensors 28702 communicate with the edge device 28704 via a hierarchical network. In these embodiments, the sensors 28702 may communicate sensor data to collection devices 206, which, in turn, may communicate the sensor data to edge device 28704 via a wired or wireless communication link. The hierarchical network may be deployed where the area being monitored is rather larger (e.g., over 40,000 sq. ft.) and power supplies are abundant, such as in a factory, a power plant, a food inspection facility, an indoor grow facility, and the like. While a hierarchal network is shown, the sensor kits 30000 of FIG. 154 may include alternative network topologies, such as a mesh topology or a star topology (e.g., sensors 28702 communicate to the edge device directly).

[1641] In example embodiments, the edge device 28704 may be configured to perform one or more AI-related tasks prior to transmission via the satellite uplink. In some of these embodiments, the edge device 28704 may be configured to determine whether there are likely no issues relating to any of the components and/or the industrial setting 28720 based on the sensor data and one or more machine-learned models. In example embodiments, the edge device 28704 may receive the sensor data from the various sensors and may generate one or more feature vectors based thereon. The feature vectors may include sensor data from a single sensor 28702, a subset of sensors 28702, or all of the sensors 28702 of the sensor kit 30000. In scenarios where a single sensor or a subset of sensors 28702 are included in the feature vector, the machine-learned model may be trained to identify one or more issues relating to an industrial component or the industrial setting 28720, but may not be sufficient to fully deem the entire setting as likely safe/free from issues. Additionally or alternatively, the feature vectors may correspond to a single snapshot in time (e.g., all sensor data in the feature vector corresponds to the same sampling event) or over a period of time (sensor data samples from a most recent sampling event and sensor data samples from previous sampling events). In embodiments where the feature vectors define sensor data from a single snapshot, the machine-learned models may be trained to identify potential issues without any temporal context. In embodiments where the feature vectors define sensor data over a period of time, the machine-learned models may be trained to identify potential issues with the context of what the sensor(s) 28702 was/were reporting previously. In these embodiments, the edge device 28704 may maintain a cache of sensor data that is sampled over a predetermined time (e.g., previous hour, previous day, previous N days), such that the cache is cleared out in a first-in-first-out manner. In these embodiments, the edge device 28704 may retrieve the previous sensor data samples from the cache to use to generate feature vectors that have data samples spanning a period of time.

[1642] In example embodiments, the edge device 28704 may feed the one or more feature vectors into one or more respective machine-learned models. A respective model may output a prediction

or classification relating to an industrial component and/or the industrial setting 28720, and a confidence score relating to the prediction or classification. In example embodiments, the edge device 28704 may make determinations relating to the manner by which sensor data is transmitted to the backend system 28750 and/or stored at the edge device. For instance, in example embodiments, the edge device 28704 may compress sensor data based on the prediction or classification. In some of these embodiments, the edge device 28704 may compress sensor data when there are no likely issues across the entire industrial setting 28720 and individual components of the industrial setting 28720. For example, if the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence (e.g., the confidence score is greater than .98), the edge device 28704 may compress the sensor data. Alternatively, in the scenario where the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence, the edge device 28704 may forego transmission but may store the sensor data at the edge device 28704 for a predefined period of time (e.g., one year). In scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may transmit the sensor data without compressing the sensor data or using a lossless compression codec. In this way, the amount of bandwidth that is transmitted via the cellular tower may be reduced, as the majority of the time the sensor data will be compressed or not transmitted.

[1643] In example embodiments, the edge device 28704 may apply one or more rules to determine whether a triggering condition exists. In example embodiments, the one or more rules may be tailored to identify potentially dangerous and/or emergency situations. In these embodiments, the edge device 28704 may trigger one or more notifications or alarms when a triggering condition exists. Additionally or alternatively, the edge device 28704 may transmit the sensor data without any compression when a triggering condition exists.

[1644] FIG. 155 illustrates an example configuration of a sensor kit 30100 for installation in an agricultural setting 30120 according to some embodiments of the present disclosure. In the example of FIG. 155, the sensor kit 30100 is configured for installation in an indoor agricultural setting 30120 that may include, but is not limited to, a control system 30122, an HVAC system 30124, a lighting system 30126, a power system 30128, and/or an irrigation system 30130. In this example, various features and components of the agricultural setting include components that are monitored by a set of sensors 28702. In example embodiments, the sensors 28702 capture instances of sensor data and provide the respective instances of sensor data to an edge device 28704. In the example embodiments of FIG. 155 the sensor kit 30100 includes a set of collection devices 206 that route sensor data from the sensors 28702 to the edge device 28704. Sensor kits 30100 for deployment in agricultural settings may have different sensor kit network topologies as well. For instance, in facilities not having more than two or three rooms being monitored, the sensor kit network may be a mesh or star network, depending on the distances between the edge device 28704 and the furthest potential sensor location. For example, if the distance between the edge device 28704 and the furthest potential sensor location is greater than 150 meters, then the sensor kit network may be configured as a mesh network. In the embodiments of FIG. 155, the edge device

28704 transmits the sensor data to the backend system 28750 directly. In these embodiments, the edge device 28704 includes a cellular communication device that communicates with a cellular tower 29910 of a preset cellular provider via a preconfigured cellular connection to a cellular tower 29910. In other embodiments of the disclosure, the edge device 28704 transmits the sensor data to the backend system 28750 via a gateway device (e.g., gateway device 30006) that includes a cellular communication device that communicates with a cellular tower 29910 of a preset cellular provider.

[1645] In example embodiments, a sensor kit 30100 may include any suitable combination of light sensors 30102, weight sensors 30104, temperature sensors 30106, CO2 sensors 30108, humidity sensors 30110, fan speed sensors 30112, and/or audio/visual (AV) sensors 30114 (e.g., cameras). Sensor kits 30100 may be arranged with additional or alternative sensors 28702. In example embodiments, the sensor data collected by the edge device 28704 may include ambient light measurements indicating an amount of ambient light detected in the area of a light sensor 30102. In example embodiments, the sensor data collected by the edge device 28704 may include a weight or mass measurements indicating a weight or mass of an object (e.g., a pot or tray containing one or more plants) that is resting upon a weight sensor 30104. In example embodiments, the sensor data collected by the edge device 28704 may include temperature measurements indicating an ambient temperature in the vicinity of a temperature sensor 30106. In example embodiments, the sensor data collected by the edge device 28704 may include humidity measurements indicating an ambient humidity in the vicinity of a humidity sensor 30110 or moisture measurements indicating a relative amount of moisture in a medium (e.g., soil) monitored by a humidity sensor 30110. In example embodiments, the sensor data collected by the edge device 28704 may include CO2 measurements indicating ambient levels of CO2 in the vicinity of a CO2 sensor 30108. In example embodiments, the sensor data collected by the edge device 28704 may include temperature measurements indicating an ambient temperature in the vicinity of a temperature sensor 30106. In example embodiments, the sensor data collected by the edge device 28704 may include fan speed measurements indicating a measured speed of a fan (e.g., a fan of an HVAC system 30124) as measured by a fan speed sensor 30112. In example embodiments, the sensor data collected by the edge device 28704 may include video signals captured by an AV sensor 30116. The sensor data captured by sensors 28702 and collected by the edge device 28704 may include additional or alternative types of sensor data without departing from the scope of the disclosure.

[1646] In example embodiments, the edge device 28704 is configured to perform one or more edge operations on the sensor data. For example, the edge device 28704 may pre-process the received sensor data. In example embodiments, the edge device 28704 may predict or classify potential issues with one or more components of the HVAC system 30124, lighting system 30126, power system 30128, the irrigation system 30130; the plants growing in the agricultural facility; and/or the facility itself. In example embodiments, the edge device 28704 may analyze the sensor data with respect to a set of rules that define triggering conditions. In these embodiments, the edge device 28704 may trigger alarms or notifications in response to a triggering condition being met. In example embodiments, the edge device 28704 may encode, compress, and/or encrypt the sensor

data, prior to transmission to the backend system 28750. In some of these embodiments, the edge device 28704 may selectively compress the sensor data based on predictions or classifications made by the edge device 28704 and/or upon one or more triggering conditions being met.

[1647] In example embodiments, the edge device 28704 may be configured to perform one or more AI-related tasks prior to transmission via the satellite uplink. In some of these embodiments, the edge device 28704 may be configured to determine whether there are likely no issues relating to any of the components and/or the industrial setting 28720 based on the sensor data and one or more machine-learned models. In example embodiments, the edge device 28704 may receive the sensor data from the various sensors and may generate one or more feature vectors based thereon. The feature vectors may include sensor data from a single sensor 28702, a subset of sensors 28702, or all of the sensors 28702 of the sensor kit 29900. In scenarios where a single sensor or a subset of sensors 28702 are included in the feature vector, the machine-learned model may be trained to identify one or more issues relating to an industrial component or the industrial setting 28720, but may not be sufficient to fully deem the entire setting as likely safe/free from issues. Additionally or alternatively, the feature vectors may correspond to a single snapshot in time (e.g., all sensor data in the feature vector corresponds to the same sampling event) or over a period of time (sensor data samples from a most recent sampling event and sensor data samples from previous sampling events). In embodiments where the feature vectors define sensor data from a single snapshot, the machine-learned models may be trained to identify potential issues without any temporal context. In embodiments where the feature vectors define sensor data over a period of time, the machine-learned models may be trained to identify potential issues with the context of what the sensor(s) 28702 was/were reporting previously. In these embodiments, the edge device 28704 may maintain a cache of sensor data that is sampled over a predetermined time (e.g., previous hour, previous day, previous N days), such that the cache is cleared out in a first-in-first-out manner. In these embodiments, the edge device 28704 may retrieve the previous sensor data samples from the cache to use to generate feature vectors that have data samples spanning a period of time.

[1648] In example embodiments, the edge device 28704 may feed the one or more feature vectors into one or more respective machine-learned models. A respective model may output a prediction or classification relating to an industrial component and/or the industrial setting 28720, and a confidence score relating to the prediction or classification. In example embodiments, the edge device 28704 may make determinations relating to the manner by which sensor data is transmitted to the backend system 28750 and/or stored at the edge device. For instance, in example embodiments, the edge device 28704 may compress sensor data based on the prediction or classification. In some of these embodiments, the edge device 28704 may compress sensor data when there are no likely issues across the entire industrial setting 28720 and individual components of the industrial setting 28720. For example, if the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence (e.g., the confidence score is greater than .98), the edge device 28704 may compress the sensor data. Alternatively, in the scenario where the machine-learned models predict that there are likely no issues and classify that there are currently no issues with a high degree of confidence, the edge

device 28704 may forego transmission but may store the sensor data at the edge device 28704 for a predefined period of time (e.g., one year). In scenarios where a machine-learned model predicts a potential issue or classifies a current issue, the edge device 28704 may transmit the sensor data without compressing the sensor data or using a lossless compression codec. In this way, the amount of bandwidth that is transmitted via the cellular tower may be reduced, as the majority of the time the sensor data will be compressed or not transmitted.

[1649] In example embodiments, the edge device 28704 may apply one or more rules to the sensor data to determine whether a triggering condition exists. In example embodiments, the one or more rules may be tailored to identify potentially dangerous and/or emergency situations. In these embodiments, the edge device 28704 may trigger one or more notifications or alarms when a triggering condition exists. Additionally or alternatively, the edge device 28704 may transmit the sensor data without any compression when a triggering condition exists. In example embodiments, the edge device 28704 may selectively compress and/or transmit the sensor data based on the application of the one or more rules to the sensor data.

[1650] In example embodiments, the backend system 28750 may perform one or more backend operations based on received sensor data. In example embodiments, the backend system 28750 may decode/decompress/decrypt the sensor data received from respective sensor kits 30100. In example embodiments, the backend system 28750 may preprocess received sensor data. In example embodiments, the backend system 28750 may preprocess sensor data received from a respective sensor kit 30100. For example, the backend system 28750 may filter, dedupe, and/or structure the sensor data. In example embodiments, the backend system 28750 may perform one or more AI-related tasks using the sensor data. In some of these embodiments, the backend system 28750 may extract features from the sensor data, which may be used to predict or classify certain conditions or events relating to the agricultural setting. For example, the backend system 28750 may deploy models used to predict yields of a crop based on weight measurements, temperature measurements, CO2 measurements, light measurements, and/or other extracted features. In another example, the backend system 28750 may deploy models used to predict or classify mold-inducing states in a room or area of the agricultural facility based on temperature measurements, humidity measurements, video signals or images, and/or other extracted features. In example embodiments, the backend system 28750 may perform one or more analytics tasks on the sensor data and may display the results to a human user via a dashboard. In example embodiments, the backend system 28750 may receive control commands from a human user via the dashboard. For example, a human resource with sufficient login credentials may control an HVAC system 30124, a lighting system 30126, a power system 30128, and/or an irrigation system 30130 of the industrial setting 28720. In some of these embodiments, the backend system 28750 may telemetrically monitor the actions of the human user, and may train one or more machine-learned models (e.g., neural networks) on actions to take in response to displaying the analytics results to the human user. In other embodiments, the backend system 28750 may execute one or more workflows associated with the HVAC system 30124, the lighting system 30126, the power system 30128, and/or the irrigation system 30130, in order to control one or more of the systems of the agricultural setting 30120 based

on a prediction or classification made by the backend system in response to the sensor data. In example embodiments, the backend system 28750 provides one or more control commands to a control system 30122 of an agricultural setting 30120, which in turn may control the HVAC system 30124, the lighting system 30126, the power system 30128, and/or the irrigation system 30130 based on the received control commands. In example embodiments, the backend system 28750 may provide or utilize an API to provide control commands to the agricultural setting 30120.

[1651] FIG. 156 illustrates an example set of operations of a method 30200 for monitoring industrial setting 28720 using an automatically configured backend system 28750. In example embodiments, the method 30200 may be performed by the backend system 28750, the sensor kit 28700, and the dashboard module 532.

[1652] At 30202, the backend system 28750 registers the sensor kit 28700 to a respective industrial setting 28720. In example embodiments, the backend system 28750 registers a plurality of sensor kits 28700 and registers each sensor kit 28700 of the plurality of sensor kits 28700 to a respective industrial setting 28720. In example embodiments, the backend system 28750 provides an interface for specifying a type of entity or industrial setting 28720 to be monitored. In example embodiments, a user may select a set of parameters for monitoring of the respective industrial setting 28720 of the sensor kit 28700. The backend system 28750 may automatically provision a set of services and capabilities of the backend system 28750 based on the selected parameters.

[1653] At 30204, the backend system 28750 configures the sensor kit 28700 to monitor physical characteristics of the respective industrial setting 28720 to which the sensor kit 28700 is registered. For example, when the respective industrial setting 28720 is a natural resource extraction setting, the backend system 28750 may configure one or more of infrared sensors, ground penetrating sensors, light sensors, humidity sensors, temperature sensors, chemical sensors, fan speed sensors, rotational speed sensors, weight sensors, and camera sensors to monitor and collect sensor data relating to metrics and parameters of the natural resource extraction setting and equipment used therein.

[1654] At 30206, the sensor kit 28700 transmits instances of sensor data to the backend system 28750. In example embodiments, the sensor kit 28700 transmits the instances of sensor data to the backend system 28750 via a gateway device. The gateway device may provide a virtual container for instances of the sensor data such that only a registered owner or operator of the respective industrial setting 28720 can access the sensor data via the backend system 28750.

[1655] At 30208, the backend system 28750 processes instances of sensor data received from the sensor kit 28700. In example embodiments, the backend system 28750 includes an analytics facility and/or a machine learning facility. The analytics facility and/or the machine learning facility may be configured based on the type of the industrial setting 28720 and may process the instances of sensor data received from the sensor kit 28700. In example embodiments, the backend system 28750 updates and/or configures a distributed ledger based on the processed instances of sensor data.

[1656] At 30210, the backend system 28750 configures and populates the dashboard. In example embodiments, the backend system 28750 configures the dashboard to retrieve and display one or

more of raw sensor data provided by the sensor kit, analytical data relating to the sensor data provided by the sensor kit 28700, predictions or classifications made by the backend system 28750 based on the sensor data, and the like. In example embodiments, the backend system 28750 configures alarm limits with respect to one or more sensor types and/or conditions based on the industrial setting 28720. The backend system 28750 may define which users receive a notification when an alarm is triggered. In example embodiments, the backend system 28750 may subscribe to additional features of the backend system 28750 and/or an edge device 28704 based on the industrial setting 28720.

[1657] At 30212, the dashboard provides monitoring information to a human user. In example embodiments, the dashboard provides monitoring information to the user by displaying the monitoring information on a device, e.g., a computer terminal, a smartphone, a monitor, or any other suitable device for displaying information. The monitoring information may be provided via a graphical user interface.

[1658] FIG. 157 illustrates an exemplary manufacturing facility 30300 according to some embodiments of the present disclosure. The manufacturing facility 30300 may include a plurality of industrial machines 30302 including, by way of example, conveyor belts, assembly machines, die machines, turbines, and power systems. The manufacturing facility 30300 may further include a plurality of products 30304. The manufacturing facility may have the sensor kit 28700 installed therein, the sensor kit 28700 including the plurality of sensors 28702 and the edge device 28704. By way of example, one or more of the sensors 28702 may be installed on some or all of the industrial machines 30302 and the products 30304.

[1659] FIG. 158 illustrates a surface portion of an exemplary underwater industrial facility 30400 according to some embodiments of the present disclosure. The underwater industrial facility 30400 may include a transportation and communication platform 30402, a storage platform 30404, and a pumping platform 30406. The underwater industrial facility 30400 may have the sensor kit 28700 installed therein, the sensor kit 28700 including the plurality of sensors 28702 and the edge device 28704. By way of example, one or more of the sensors 28702 may be installed on some or all of the transportation and communication platform 30402, the storage platform 30404, and the pumping platform 30406, and on individual components and machines thereof.

[1660] FIG. 159 illustrates an exemplary indoor agricultural facility 30500 according to some embodiments of the present disclosure. The indoor agricultural facility 30500 may include a greenhouse 30502 and a plurality of wind turbines 30504. The indoor agricultural facility 30500 may have the sensor kit 28700 installed therein, the sensor kit 28700 including the plurality of sensors 28702 and the edge device 28704. By way of example, one or more of the sensors 28702 may be installed on some or all components of the greenhouse 30504 and on some or all components of the wind turbines 30504.

[1661] Referring to FIG. 160, in example embodiments, the edge device 28704 may include, link or connect to, integrate with, or be integrated into the control system 13742 and/or the data handling platform 13700 for providing control for one or more industrial entities 13736, such as controlling a machine in a factory (such as a CNC machine, additive manufacturing machine, energy system

(e.g., a generator or turbine), an assembly line, or the like), controlling a workflow (such as a production workflow, an inspection workflow, a data collection workflow, a maintenance workflow, a servicing workflow, or the like), or controlling sub-systems, systems, or operations of an entire factory or set of factories. In example embodiments, the edge device 28704 may link or connect to the control system 13742 via the network 28780. In example embodiments, the edge device 28704 may integrate with the control system 13742 via the processing device 29006. In example embodiments, the control system 13742 may integrate with the backend system 28750. Processing, computation and intelligence capabilities of the edge device 28704 may thus benefit from input from a set of control systems 13742 and may provide inputs to (including control signals for) the set of control systems 13742. Data from the sensor kit 28700 (including reporting packets, sensor kit packets, and/or other data from sensors 28702 and/or the data processing module 29020, the encoding module 29022, the quick-decision AI module 29024, the notification module 29026, the configuration module 29028, and the distributed ledger module 29030), and/or from the edge device 28704 may be represented in the set of industrial digital twins 13734. For example, an industrial digital twin 13734 may show a point cloud view of the industrial setting 28720 (which, in example embodiments, may be augmented, such as using 3D mapping, AR or VR systems) with relevant data collection elements presented in the point cloud view along with the point cloud. Many examples are available, such as highlighting (such as by color or motion) in the digital twin 13734, areas of the point cloud where systems are vibrating in a way that is out of the normal range (such as where severity units, as discussed elsewhere herein, exceed a threshold). Industrial entity digital twins 13734 may include, link or connect to, or integrate with a variety of interfaces and dashboards 13738, such as ones configured for specific workflows, roles, and users. For example, dashboards and interfaces may be configured for workers who will interact with specific machines (such as where the digital twin is used for training, workflow guidance, diagnosis of problems, and the like); for managers of operations on a factory floor (such as where a digital twin 13734 displays a layout of machines on the floor, patterns of traffic (e.g., moving assets. 13708 and workers 13712) involved in workflows, status information for workers, machines, processes, or the like (including operational status, maintenance status, inspection status, and the like), analytic information (such as indicating metrics about operations, about potential problems, or the like); for inspectors (such as where the digital twin 13734 represents areas that are indicated by data collectors 13702 to require or benefit from additional inspection (e.g., where the inspector can check off items that have already been inspected or highlight items for further inspection by interacting with them in a digital twin interface or dashboard 13738); for maintenance and service workers (such as where a digital twin 13734 highlights locations of items requiring maintenance in a schematic view and guides the service workers to the right location and/or machine, then presents (such as in a different view) information and guidance on how to undertake the service or maintenance, ranging from a checklist or workflow to a virtual, mixed or augmented reality training or guidance session that can be presented at the machine); for front office managers (such as finance professionals who can be presented financial information, such as ROI metrics, output metrics, cost metrics, and the like (including current status and predictions), legal personnel (such as where a digital twin 13734 may

present compliance information, highlight legal risks (such as safety violations or instances where status information about operations indicates a likelihood that the company may breach a contract (such as by failing to produce an output that is required by a contract) and the like), inventory managers, procurement personnel, and the like; and for executives, such as CEOs, CTOs, COOs, CIOs, CDOs, CMOs, and the like, who may interact with digital twins 13734 that represent whole factories, or sets of factories, such as to identify risks and opportunities that may involve understanding interactions of elements and/or contributions of elements involving industrial entities 13736 to overall operations of an enterprise, to its strategies, or the like. The digital twin 13734 may be updated based upon data from the sensor kit 28700 such that the digital twin 13734 is maintained in substantially real time.

[1662] In various embodiments, the interfaces and dashboards 13738 may display sensor information collected from the sensor kit 28700. Information elements from the industrial environment 13704 or about industrial setting 28720 can be presented in overlays (e.g., where metrics or symbols are presented on top of a point cloud, a photo, or a 3D representation of a unit in a 3D interface), in native form (such as where a point cloud is represented), in 3D visualizations (such as where the interface handles elements as 3D geometric elements), and the like.

[1663] Systems and methods for using wearable devices for mobile data collection within an environment for industrial IoT data collection are next described with respect to FIGS. 161 to 164. Referring first to FIG. 161, a data collection system may include one or more wearable devices configured to act as mobile data collectors within an environment for industrial IoT data collection. For example, the one or more wearable devices may transmit data to, receive data from, transmit commands to, receive commands from, be under the control of, communicate controls for, or otherwise communicate with the industrial IoT data collection, monitoring and control system 10. Methods and systems are disclosed herein for data collection using wearable devices, including a single wearable device having a single sensor for recording state-related measurements (otherwise "measurements of states" or "state measurements," as noted below) within the environment for industrial IoT data collection, a single wearable device having multiple sensors for recording state-related measurements within the environment for industrial IoT data collection, multiple wearable devices each having a single sensor for recording state-related measurements within the environment for industrial IoT data collection, and multiple wearable devices each having one or more sensors for recording state-related measurements within the environment for industrial IoT data collection. For example, a wearable device may be a wearable haptic or multi-sensor user interface for an industrial sensor data collector, with vibration, heat, electrical, and/or sound outputs, and any other suitable outputs. In another example, a wearable device may be any other suitable device, component, unit, or other computational aspect having a tangible form and which is configured or otherwise able to be used by disposing on a person within an industrial environment, regardless of the period of time of such use. For example, a wearable device may be an article of clothing or a device included within an article of clothing. In another example, a wearable device may be an accessory article or a device included within an accessory article. Examples of articles of clothing that the wearable device can be or be included within include,

without limitation, shirts, vests, jackets, pants, shorts, gloves, socks, shoes, protective outerwear, undergarments, undershirts, tank tops, and the like. Examples of accessory articles that the wearable device can be or be included within include, without limitation, hats, helmets, glasses, goggles, vision safety accessories, masks, chest bands, belts, lift support garments, antennae, wrist bands, rings, necklaces, bracelets, watches, brooches, neck straps, backpacks, front packs, arm packs, leg packs, lanyards, key rings, headphones, hearing safety accessories, earbuds, earpieces, and the like. Regardless of the particular form, a wearable device according to this disclosure includes one or more sensors for recording state-related measurements of an environment for industrial IoT data collection. For example, the one or more sensors of a wearable device described in this disclosure can measure states with respect to equipment within an industrial IoT environment or with respect to the industrial IoT environment itself. As used herein, a measurement of a state recorded using a sensor (e.g., of a wearable device or of any other suitable data collector) refers to information relating to a target of the environment for industrial IoT data collection. That is, the information directly or indirectly indicates a state of a target, or may otherwise be used to indicate a state of a target. For example, the information may indirectly indicate a state of a target where it is processed or otherwise used to identify or determine the state of the target. As used herein, the recording of a measurement using a sensor (e.g., of a wearable device or of any other suitable data collector) refers to the use of the sensor in making the measurement available for further processing. For example, recording a measurement using a sensor may refer to one or more of generating data indicative of the measurement, transmitting a signal indicative of the measurement, or otherwise obtaining values for the measurement.

[1664] A number of wearable devices 14000 are located within the environment for industrial IoT data collection. In some scenarios, the wearable devices 14000 may be wearable devices issued by an operator of the environment for industrial IoT data collection. Alternatively, the wearable devices 14000 may be wearable devices owned by workers selected to perform tasks within the environment for industrial IoT data collection. As shown in FIG. 161, the wearable devices 14000 may include any combination of a single wearable device with a single sensor 14002, a single wearable device with multiple sensors 14004, a combination of wearable devices each with a single sensor 14006, and a combination of wearable devices each with one or more sensors 14008. However, in example embodiments, the wearable devices 14000 may include different wearable devices. For example, in example embodiments, the wearable devices 14000 may omit the combination of wearable devices each with a single sensor 14006 and/or the combination of wearable devices each with one or more sensors 14008. For example, the wearable devices 14000 may be limited to individual wearable devices rather than combinations of wearable devices that offer combined, improved or otherwise different functionality when compared to each of the constituent wearable devices taken individually. In another example, in example embodiments, the wearable devices 14000 may omit the single wearable device with the single sensor 14002 and/or the single wearable device with multiple sensors 14004. For example, the wearable devices 14000 may be limited to combinations of wearable devices rather than individual devices (e.g., where specific combinations of the wearable devices are identified as being valuable in particular contexts

or otherwise for recording particular state-related measurements within the environment for industrial IoT data collection). Communications and other transfers of data between the wearable devices 14000 and the devices that receive the output from the wearable devices, or otherwise between the sensors within the wearable devices 14000 and a device that receives the output of those sensors, may be wireless or wired and may include such standard communication technologies as 802.11 and 900 MHz wireless systems, Ethernet, USB, firewire, and so on

[1665] In example embodiments, different wearable devices 14000 may be configured to record certain types of state-related measurements of some or all of the targets (e.g., devices or equipment) within the environment for industrial IoT data collection. For example, some of the wearable devices 14000 may be configured to record state-related measurements of targets based on vibrations measured with respect to some or all of the targets. A vibration measured with respect to a target may refer to, without limitation, a frequency at which all or a portion of the target vibrates, a waveform derived from a vibration envelope associated with the target, vibration level changes, and the like. In another example, some of the wearable devices 14000 may be configured to record state-related measurements of targets based on temperatures measured with respect to some or all of the targets. A temperature measured with respect to a target may refer to, without limitation, an internal or external temperature of all or a portion of the target, an operating temperature of the target, a temperature measured within an area around the target, and the like. In another example, some of the wearable devices 14000 may be configured to record state-related measurements of targets based on electrical or magnetic outputs measured with respect to some or all of the targets. An electrical or magnetic output measured with respect to a target may refer to, without limitation, a level or change in an electromagnetic field associated with the target, an amount of electricity or magnetic quality output from the target or otherwise emitted by the target, and the like. In another example, some of the wearable devices 14000 may be configured to record state-related measurements of targets based on sound outputs measured with respect to some or all of the targets. A sound output measured with respect to a target may refer to, without limitation, an audible or inaudible frequency corresponding to a sound wave generated by or in connection with the target, a sound wave emitted by a change in operation of the target, and the like. In another example, some of the wearable devices 14000 may be configured to record state-related measurements of targets based on outputs other than vibrations, temperatures, electrical or magnetic, or sound, as measured with respect to some or all of the targets.

[1666] Alternatively, or additionally, different wearable devices 14000 may be configured to record some or all state-related measurements of certain types of the targets within the environment for industrial IoT data collection. For example, some of the wearable devices 14000 may be configured to record some or all state-related measurements from agitators (e.g., turbine agitators), airframe control surface vibration devices, catalytic reactors, compressors and the like. In another example, some of the wearable devices 14000 may be configured to record some or all state-related measurements from conveyors and lifters, disposal systems, drive trains, fans, irrigation systems, motors, and the like. In another example, some of the wearable devices 14000 may be configured to record some or all state-related measurements from pipelines, electric powertrains, production

platforms, pumps (e.g., water pumps), robotic assembly systems, thermic heating systems, tracks, transmission systems, turbines, and the like. In example embodiments, the wearable devices 14000 may be configured to record some or all state-related measurements of certain types of industrial environments. For example, an industrial environment having targets with states measured using the wearable devices 14000 may include, but is not limited to, a manufacturing environment, a fossil fuel energy production environment, an aerospace environment, a mining environment, a construction environment, a ship environment, a shipping environment, a submarine environment, a wind energy production environment, a hydroelectric energy production environment, a nuclear energy production environment, an oil drilling environment, an oil pipeline environment, any other suitable energy product environment, any other suitable energy routing or transmission environment, any other suitable industrial environment, a factory, an airplane or other aircraft, a distribution environment, an energy source extraction environment, an offshore exploration site, an underwater exploration site, an assembly line, a warehouse, a power generation environment, a hazardous waste environment, and the like.

[1667] The combination of wearable devices each with a single sensor 14006 and/or the combination of wearable devices each with one or more sensors 14008 may represent a combination of wearable devices selected for use together within the environment for industrial IoT data collection. For example, the combination of wearable devices each with a single sensor 14006 and/or the combination of wearable devices each with one or more of the sensors 14008 may represent all or a portion of an industrial uniform to be worn by a worker performing one or more tasks within the environment for industrial IoT data collection. For example, the combination of wearable devices each with the single sensor 14006 and/or the combination of wearable devices each with one or more of the sensors 14008 may include one of each of a number of wearable devices to be worn by the user (e.g., one hat, one shirt, one pair of pants, one pair of shoes, one vest, one necklace, one bracelet, one backpack, or more or fewer wearable devices). Embodiments of this disclosure may contemplate industrial uniforms as including other possible combinations of the wearable devices as the combination of wearable devices each with the single sensor 14006 and/or the combination of wearable devices each with one or more of the sensors 14008.

[1668] In example embodiments, the combined use of multiple sensors, either as the combination of wearable devices each with the single sensor 14006 and/or as the combination of wearable devices each with one or more of the sensors 14008, may introduce extended or additional functionality for industrial IoT data collection. Thus, in some of those embodiments, an industrial uniform may include functionality beyond what is provided by the individual sensors that are integrated within the industrial uniform. For example, the output of wearable devices with sensors for recording state-related measurements of the same target may be pre-processed by a central processing software or hardware aspect integrated within or otherwise corresponding to the industrial uniform (e.g., a collective processing mind, as described below). For example, the central processing software or hardware aspect integrated within or otherwise corresponding to the industrial uniform may process the output of multiple wearable devices to determine whether the output is the same for a particular observed measurement of a target. Where one of those outputs

is more than a threshold deviation from the other outputs, that deviated output may be discarded. For example, the discarded output may represent output produced using a sensor that suffered from interference or other issues while recording the state-related measurement of the target. In another example, the central processing software or hardware aspect integrated within or otherwise corresponding to the industrial uniform may process different types of output (e.g., recorded based on different targets or different state-related measurement types, for example, vibrational versus temperature) of multiple wearable devices to determine or identify a state of the target. For example, it may be the case that a state is indicated by a combination of outputs. In such a scenario, a first output from a first wearable device can be combined or otherwise processed along with a second output from a second wearable device to determine or identify the state of the target. Different combinations of wearable devices may be identified as different industrial uniforms, in which each of the industrial uniforms may have the same or different capabilities with respect to recording types of state-related measurements of targets. In yet another example, the integration of multiple wearable devices within an industrial uniform allows for the concurrent or substantially concurrent processing of state-related measurements recorded using those wearable devices.

[1669] The state-related measurements using the wearable devices 14000 may be made available over a network 14010 (e.g., without the need for external networks). The network 14010 may be a MANET (e.g., the MANET 20 shown in Figure 2 or any other suitable MANET), the Internet (e.g., the Internet 110 shown in Figure 3 or any other suitable Internet), or any other suitable type of network, or any combination thereof. For example, the network 14010 may be used to receive state-related measurements recorded using the wearable devices 14000. The network 14010 may then be used to transmit some or all of those received state-related measurements to other components of the data collection system 102. For example, the network 14010 may be used to transmit some or all of the received state-related measurements to a data pool 14012 (e.g., the data pool 60 shown in Figure 2 or any other suitable data pool) for storage of those received state-related measurements. In another example, the network 14010 may be used to transmit some or all of the received state-related measurements to one or more servers 14014 corresponding to the environment for industrial IoT data collection. The servers 14014 may include one or more hardware or software server aspects. For example, the servers 14014 to which the received state-related measurements are transmitted may include intelligent systems 14016 that process the received state-related measurements. The intelligent systems 14016 may process the received state-related measurements in any suitable manner, including using artificial intelligence processes, machine learning processes, and/or other cognitive processes to identify information within or otherwise associated with the received state-related measurements. In example embodiments, after processing the received state-related measurements, the servers 14014 to which the received state-related measurements are transmitted may transmit the processed information or data indicative of the processed information to other systems (e.g., for storage or analysis). The data indicative of the processed information from the servers 14014 may include, for example, output or other results of the artificial intelligence processes, machine learning processes, and/or other cognitive processes.

[1670] In example embodiments, some or all of the wearable devices 14000 may include intelligent systems 14018 for processing the state-related measurements recorded using those wearable devices 14000 before transmitting those recorded state-related measurements (e.g., over the network 14010) or any other suitable communication mechanism. For example, some or all of the wearable devices 14000 may integrate artificial intelligence processes, machine learning processes, and/or other cognitive processes for analyzing the state-related measurements recorded thereby. The processing by the intelligent systems 14018 of the wearable devices 14000 may be or be represented within a pre-processing step of the industrial IoT data collection, monitoring and control system 10. For example, the pre-processing may be selectively performed by certain types of the wearable devices 14000 to pre-process the recorded state-related measurements, for example, to identify redundant information, irrelevant information, or insignificant information. In another example, the pre-processing may be automated for certain types of the wearable devices 14000 to pre-process the recorded state-related measurements, for example, to identify redundant information, irrelevant information, or insignificant information. In another example, the pre-processing may be selectively performed for certain types of state-related measurements recorded by any of the wearable devices 14000 to pre-process the recorded state-related measurements, for example, to identify redundant information, irrelevant information, or insignificant information. In another example, the pre-processing may be automated for certain types of state-related measurements recorded by any of the wearable devices 14000 to pre-process the recorded state-related measurements, for example, to identify redundant information, irrelevant information, or insignificant information.

[1671] In example embodiments, some or all of the wearable devices 14000 may include sensor fusion functionality. For example, the sensor fusion functionality may be embodied as the on-device sensor fusion 80. For example, state-related measurements recorded using multiple analog sensors of one or more of the wearable devices 14000 (e.g., the multiple analog sensors 82 shown in Figure 4 or any other suitable sensor) may be locally or remotely processed (e.g., using artificial intelligence processes, machine learning processes, and/or other cognitive processes), which may be embodied within the wearable devices 14000 themselves, within the servers 14014, within both, or within any other suitable hardware or software. For example, the output of the sensors integrated within the wearable devices 14000 may be provided directly to the on-device sensor fusion aspect 80. The sensor fusion functionality may be embodied by a pre-processing step that is performed prior to the artificial intelligence processes, machine learning processes, and/or other cognitive processes. In example embodiments, the sensor fusion functionality may be performed using a MUX. For example, each of the single wearable devices with multiple sensors 14004 may include its own MUX for combining state-related measurements recorded using different individual sensors of those multiple sensors. In another example, some or all of the individual wearable devices within the combination of wearable devices each with one or more sensors 14008 may include its own MUX for combining state-related measurements recorded using different individual sensors of those multiple sensors. In some such embodiments, the MUX may be internal

to those wearable devices. In some such embodiments, the MUX may be external to those wearable devices.

[1300] In example embodiments, the wearable devices 14000 may be controlled by or otherwise used in connection within a host processing system 112 shown in Figure 6 (or any other suitable host system). The host processing system 112 may be locally accessible over the network 14010. Alternatively, the host processing system 112 may be remote (e.g., embodied in a cloud computing system), may be accessible using one or more network infrastructure elements (e.g., access points, switches, routers, servers, gateways, bridges, connectors, physical interfaces and the like), and/or may use one or more network protocols (e.g., IP-based protocols, TCP/IP, UDP, HTTP, BLUETOOTH™, BLUETOOTH LOW ENERGY™, cellular protocols, LTE, 2G, 3G, 4G, 5G, CDMA, TDSM, packet-based protocols, streaming protocols, file transfer protocols, broadcast protocols, multi-cast protocols, unicast protocols, and the like). In example embodiments, the state-related measurements recorded using the wearable devices 14000 may be processed using a network coding system or method, which may be embodied internally or externally with respect to the host processing system 112. For example, the network coding system can process the measurements recorded using the wearable devices 14000 based on the availability of networks for communicating those recorded state-related measurements, based on the availability of bandwidth and spectrum for communicating those recorded state-related measurements, based on other network characteristics, or based on some combination thereof.

[1672] In example embodiments, the state-related measurements recorded using the wearable devices 14000 may be pulled from the wearable devices 14000 by an upstream device (e.g., a client device or other software or hardware aspect used to review, analyze, or otherwise view the state-related measurements). For example, the wearable devices 14000 may not actively transmit the state-related measurements that are received (e.g., at the servers 14014, the data pool 14012, or any other suitable hardware or software component that receives the state-related measurements recorded using the wearable devices 14000). Rather, the transmission of the state-related measurements from the wearable devices 14000 may be caused by commands received at the wearable devices 14000 (e.g., from servers 14014 or from other hardware or software of the data collection system 102). For example, a data collector, which may be fixed within a particular location of the environment or which may be mobile with respect to the environment, may be configured to pull state-related measurements recorded by various wearable devices 14000. For example, the wearable devices 14000 may continuously, periodically, or otherwise at multiple times record state-related measurements within the environment for industrial IoT data collection. The data collector may, at fixed intervals, at random times, or otherwise, transmit one or more commands to some or all of the wearable devices 14000 (e.g., to pull some or all of the state-related measurements recorded by those wearable devices 14000 since the last time state-related measurements were pulled therefrom). Alternatively, the data collector may, at those fixed intervals, at those random times, or otherwise, transmit the one or more commands to a collective processing mind 14020 associated with the wearable devices 14000. For example, the collective processing mind 14020 may be or include a hub for receiving the state-related measurements

recorded using some or all of the wearable devices 14000. In another example, the commands, when processed using individual wearable devices 14000 or by the collective processing mind 14020 of the wearable devices 14000, cause the recorded state-related measurements or data representative thereof to be transmitted from the wearable devices 14000. For example, the collective processing mind 14020 may be configured to pull the state-related measurements from some or all of the wearable devices 14000 (e.g., at fixed intervals, at random times, or otherwise). The collective processing mind 14020 may then transmit the state-related measurements pulled from the wearable devices 14000 (e.g., to the servers 14014, the data pool 14012, or the other hardware or software component selected or otherwise configured to receive the state-related measurements).

[1673] In example embodiments, the state-related measurements recorded using the wearable devices 14000 may be transmitted from the wearable devices 14000 responsive to requests for those state-related measurements. For example, the collective processing mind 14020 may, at fixed intervals, at random times, or otherwise, transmit a request for recorded state-related measurements to some or all of the wearable devices 14000. The processors of some or all of the wearable devices 14000 to which the request is sent may process the request to determine which state-related measurements to transmit. For example, data indicative of a time of a most recent request for recorded state-related measurements may be accessed by those processors. The processors may then compare that time to a time at which the new request is received from the collective processing mind 14020. The processors may then query a data store for state-related measurements recorded between the two times. The processors may then transmit those state-related measurements in response to the request. In another example, the processors may identify a most recent set of state-related measurements recorded using the corresponding wearable devices 14000 and transmit those state-related measurements in response to the request. In another example, data collectors within the data collection system 10 may transmit the request directly to the wearable devices 14000. In yet another example, the data collectors may transmit the request to the collective processing mind 14020. The collective processing mind 14020 may process the request to determine select individual wearable devices 14000 which were used to record the requested state-related measurements. The collective processing mind 14020 may then transmit certain state-related measurements in response to the request by, for example, querying a storage for some or all of the state-related measurements recorded using those select individual wearable devices 14000. Alternatively, the collective processing mind 14020 may process the request to determine which of the state-related measurements recorded by some or all of the wearable devices 14000 to transmit in response to the request (e.g., based on a time of the request). For example, the collective processing mind 14020 can compare the time of the request to a time of a most recent request for recorded state-related measurements. The collective processing mind 14020 can then retrieve the state-related measurements recorded in between those times and transmit the retrieved state-related measurements in response to the request.

[1674] In example embodiments, the state-related measurements may be pushed from the wearable devices 14000 to an upstream device (e.g., a client device or other software or hardware aspect

used to review, analyze, or otherwise view the state-related measurements). For example, the wearable devices 14000 may actively transmit the state-related measurements that are received (e.g., to the servers 14014, the data pool 14012, or any other suitable hardware or software component that receives the state-related measurements recorded using the wearable devices 14000) without such receiving hardware or software component requesting those state-related measurements or otherwise causing the wearable device to transmit those state-related measurements based on a command. For example, some or all of the wearable devices 14000 may transmit state-related measurements on a fixed interval, at random times, immediately upon the recording of those state-related measurements, some amount of time after recording those measurements, upon a determination that a threshold number of state-related measurements have been recorded, or at other suitable times. In some such embodiments, the wearable devices 14000, either by themselves or using the collective processing mind 14020, may push the recorded state-related measurements in response to detecting a near proximity of a data collection router 14014.

[1675] For example, referring next to Figure 162, the collective processing mind 14020 may include a detector 14022 configured to detect a near proximity of a target 14024 (e.g., one of the devices 13006 shown in Figure 134 or any other suitable target) with respect to one or more of the wearable devices 14000. For example, upon such a detection, the collective processing mind 14020 may send a signal to the one or more of the wearable devices 14000 to record and transmit state-related measurements of receipt at a data collection router 14026. Alternatively, upon such a detection, the collective processing mind 14020 may query a data store to retrieve state-related measurements and then transmit those state-related measurements of receipt at the data collection router 14026. In either case, the data collection router 14026 forwards the received state-related measurements to the servers 14014, the data pool 14012, or any other suitable hardware or software component. In another example, upon such a detection, the collective processing mind 14020 may send the signal directly to the servers 14014, the data pool 14012, or the other hardware or software component, for example, to bypass the data collection router 14026 or where the data collection router 14026 is omitted.

[1676] Referring next to Figure 163, in example embodiments, the collective processing mind 14020 may be omitted. In some of these embodiments, the wearable devices 14000 may detect the near proximity of the target 14024. Upon such detection, the wearable devices 14000 may record state-related measurements of the target 14024 (e.g., vibrations, temperature, electrical or magnetic output, sound output, or the like). The recorded state-related measurements can be transmitted over the network 14010 (e.g., to the data pool 14012, the servers 14014, or any other suitable hardware or software component). Alternatively, the recorded state-related measurements can be transmitted to the data collection router 14026, for example, where the network 14010 is unavailable or where the data collection router 14026 is configured to receive and/or pre-process the recorded state-related measurements from the wearable devices 14000. The data collection router 14026 may be one of a number of data collection routers 14026 located throughout the environment for industrial IoT data collection. For example, the data collection router 14026 may be the data collection router 14026 configured to transmit state-related measurements specifically recorded for the target 14024.

[1677] Referring next to Figure 164, various aspects of functionality of intelligent systems 14028 used to process output of the wearable devices 14000 are disclosed. In example embodiments, the intelligent systems 14028 include a cognitive learning module 14030, an artificial intelligence module 14032, and a machine learning module 14034. The intelligent systems 14028 may include additional or fewer modules. The intelligent systems 14028 may, for example, be the intelligent systems 14018 or the intelligent systems 14016 shown in Figure 161 or other intelligent systems. Although shown as separate modules, in example embodiments, there may be an overlap between some or all of the cognitive learning module 14030, the artificial intelligence module 14032, and the machine learning module 14034. For example, the artificial intelligence module 14032 may include the machine learning module 14034. In another example, the cognitive learning module 14030 may include the artificial intelligence module 14032 (and, in example embodiments, therefore, the machine learning module 14034). The wearable devices 14000 may include any number of wearable devices. For example, as shown, the wearable devices 14000 include a first wearable device 14000A, a second wearable device 14000B, and an Nth wearable device 14000N, where N is a number greater than two. The intelligent systems 14028 receives the output of the wearable devices 14000A, 14000B, ... 14000N. In particular, one or more of the modules 14030, 14032, and 14034 of the intelligent systems 14028 receives data generated by and output from one or more of the wearable devices 14000A, 14000B, ... 14000N. The output from the wearable devices 14000A, 14000B, ... 14000N may, for example, include state-related measurements recorded using the wearable devices 14000A, 14000B, ... 14000N (e.g., state-related measurements of equipment within an environment for industrial IoT data collection). In example embodiments, the output from the wearable devices 14000A, 14000B, ... 14000N may be processed by all three of the modules 14030, 14032, and 14034 of the intelligent systems 14028. In example embodiments, the output from the wearable devices 14000A, 14000B, ... 14000N may be processed by only one of the modules 14030, 14032, and 14034 of the intelligent systems 14028. For example, the particular one of the modules 14030, 14032, and 14034 of the intelligent systems 14028 to use to process the output from the wearable devices 14000A, 14000B, ... 14000N may be selected based on the wearable device used to generate that output, the equipment measured in generating that output, the values of the output, other selection criteria, and the like.

[1678] A knowledge base 14036 may be updated based on output from the intelligent systems 14028. The knowledge base 14036 represents a library or other set or collection of knowledge related to the environment of the industrial IoT data collection, including equipment within that environment, tasks performed within that environment, personnel having the skill to perform tasks within that environment, and the like. The intelligent systems 14028 can process the state-related measurements recorded using the wearable devices 14000A, 14000B, ... 14000N to facilitate knowledge gathering for expanding the knowledge base 14036. For example, the modules 14030, 14032, and 14034 of the intelligent systems 14028 can process those state-related measurements against existing knowledge within the knowledge base 14036 to update or otherwise modify information within the knowledge base 14036. The intelligent systems 14028 may use intelligence and machine learning capabilities (e.g., of the machine learning module 14034 or as described

elsewhere in this disclosure) to process state-related measurements and related information based on detected conditions (e.g., conditions informed by the wearable devices 14000 and/or provided as training data) and/or state information (e.g., state information determined by a machine state recognition system that may determine a state, for example, information relating to an operational state, an environmental state, a state within a known process or workflow, a state involving a fault or diagnostic condition, and the like). This may include optimization of input selection and configuration based on learning feedback from the learning feedback system, which may include providing training data (e.g., from a host processing system or from other data collection systems either directly or from the host processing system) and may include providing feedback metrics (e.g., success metrics calculated within an analytic system of the host processing system). Examples of host processing systems, learning feedback systems, data collection systems, and analytic systems are described elsewhere in this disclosure. Thus, the intelligent systems 14028 can be used to update workflows of tasks assigned and performed within the industrial IoT environment based on output from the wearable devices 14000A, 14000B, ... 14000N.

[1679] In example embodiments, the intelligent systems 14028, either within one of the modules 14030, 14032, and 14034 or otherwise, may include other intelligence or machine learning aspects. For example, the intelligent systems 14028 may include one or more of a you only look once (YOLO) neural network, a YOLO convolutional neural network (CNN), a set of neural networks configured to operate on or from a FPGA, a set of neural networks configured to operate on or from a FPGA and graphics processing unit (GPU) hybrid component, a user configurable series and parallel flow for a hybrid neural network (e.g., configuring series and/or parallel flows between neural networks as outputs which can be communicated between such neural networks), a machine learning system for automatically configuring a topology or workflow for a set of hybrid neural networks (e.g., series, parallel, data flows, etc.) based on a training data set which may or may not use manual configurations (e.g., by a human user), a deep learning system for automatically configuring a topology or workflow for a set of hybrid neural networks (e.g., series, parallel, data flows, etc.) based on a training data set of outcomes from industrial IoT processes (e.g., maintenance, repair, service, prediction of faults, optimization of operation of a machine, system of facility, etc.), or other intelligence or machine learning aspects.

[1680] Thus, in example embodiments, the output of the wearable devices 14000 may be processed using the intelligent systems 14028 to add to, remove from, or otherwise modify the knowledge base 14036. For example, the knowledge base 14036 may reflect information to use to perform one or more tasks within the industrial environment in which the targets are located and in which the wearable devices 14000 are used. The output from the wearable devices 14000 can thus be used to increase knowledge as to the nature of issues that arise with respect to the industrial environment, for example, by describing information about the target from which measurements were recorded, a time and/or date at which the measurements were recorded, pre-existing state or other condition information about the target, information about the time required to resolve an issue with respect to a target, information about how to resolve an issue with respect to a target, information indicating an amount of downtime to the target and to other aspects of the respective

industrial environment resulting from resolving the issue, an indication of whether the issue should be resolved now or later (or not at all), and the like. The intelligent systems 14028 may process that output to update existing training data. For example, the existing training data can be used to update the machine learning, artificial intelligence, and/or other cognitive functionality for identifying states of targets based on the output of the wearable devices 14000.

[1681] For example, the knowledge base 14036 may include a series of databases or other tables or graphs arranged hierarchically based on the target or the area of the industrial environment that includes the target. For example, a first layer of the knowledge base 14036 may refer to the industrial environment (e.g., a power plant, a manufacturing facility, a mining facility, etc.). A second layer of the knowledge base 14036 may refer to zones within the industrial environment (e.g., zone 1, zone 2, etc., or named zones, as the case may be). A third layer of the knowledge base 14036 may refer to targets within those zones (e.g., within a first zone of a power plant including electrical equipment, this could include alternators, circuit breakers, transformers, batteries, exciters, etc., and, within a second zone of a power plant including a turbine, a generator, a generator magnet, etc.). The knowledge base 14036 may be updated based on output of the intelligent systems 14028, by manual user data entry, or both. For example, a worker within a power plant may be given one or more wearable devices (e.g., the wearable devices 14000). In approaching a turbine, one of the wearable devices 14000 having a sensor for recording vibrational measurements may determine that the turbine is vibrating at a particular rate. The output of the wearable device is processed by the intelligent systems 14028, such as by comparing that output against the set of known data for the turbine. For example, the intelligent systems 14028 can query data from the knowledge base 14036 indicating historical measurements recorded with respect to the vibrations of that turbine within that particular power plant. The intelligent systems 14028 can then determine whether the new output from the wearable device is consistent with the data within the knowledge base 14036 or is deviant therefrom. In the event the new output deviates from the data within the knowledge base, the intelligent systems 14028 can update the data within that portion of the knowledge base 14036 to reflect the new output. Alternatively, the updating of the knowledge base 14036 may be delayed, for example, until after a threshold number of deviant output measurements are recorded, so as to prevent misrepresentative output from being used to modify the operational understanding of the turbine.

[1682] Disclosed herein are systems for data collection in an industrial environment with wearable device integration. As used herein, wearable device integration refers to using wearable devices for specific or general purposes. For example, wearable device integration as described with respect to the functionality or configuration of a system refers to the use by that system of the wearable devices 14000 and/or the hardware and/or software used in connection with the wearable devices 14000 for data collection within an industrial IoT environment, for example, as shown in FIGS. 161 to 164. Such wearable device integration refers to the use of one or more of the wearable devices 14000. For example, a system disclosed herein as including wearable device integration may include integration of one or more of a shirt, vest, jacket, pair of pants, pair of shorts, glove, sock, shoe, protective outerwear, undergarment, undershirt, tank top, hat, helmet, glasses, goggles,

vision safety accessory, mask, chest band, belt, lift support garment, antenna, wrist band, ring, necklace, bracelets, watch, brooch, neck strap, backpack, front pack, arm pack, leg pack, lanyard, key ring, headphones, hearing safety accessory, earbuds, or earpiece, or of other types of wearable devices or articles (e.g., articles of clothing and/or accessory articles) including such other types of wearable devices.

[1683] In example embodiments, a mobile data collector swarm 14038 includes a number of mobile robots and/or mobile vehicles. The mobile robots and/or mobile vehicles of the swarm 14038 may be mobile robots and/or mobile vehicles native to the industrial IoT environment or mobile robots and/or mobile vehicles brought into the industrial IoT environment from a different location. As shown in Figure 165, the swarm 14038 may include different types of mobile robots and/or mobile vehicles, including a mobile robot with one or more mobile data collectors integrated therein 14040, a mobile vehicle with one or more mobile data collectors integrated therein 14042, a mobile robot with one or more mobile data collectors coupled thereto 14044, and a mobile vehicle with one or more mobile data collectors coupled thereto 14046. In example embodiments, a mobile data collector is integrated within a mobile robot or mobile vehicle when removal of the mobile data collector from the mobile robot or mobile vehicle during the typical operation of the mobile robot or mobile vehicle would result in disruption to the principle operation of the mobile robot or mobile vehicle. In example embodiments, a mobile data collector is coupled to a mobile robot or mobile vehicle when the mobile data collector is able to be removed or otherwise uncoupled from the mobile robot or mobile vehicle without material disruption to the principle operation of the mobile robot or mobile vehicle.

[1684] The mobile robots and mobile vehicles of the mobile data collector swarm 14038 collect data from targets 14048 (e.g., the targets 12002 shown in Figure 118, or any other suitable target). In example embodiments, data collected by the mobile data collectors from the targets 14048 can be stored in a data pool 14050 (e.g., the data pool 14012 shown in FIG. 161, or any other suitable data pool). For example, the targets 14048 may be or include one or more of machines, pipelines, equipment, installations, tools, vehicles, turbines, speakers, lasers, automatons, computer equipment, industrial equipment, switches, and the like.

[1685] Different mobile robots and/or mobile vehicles of the swarm 14038 may be configured to record certain types of state-related measurements of some or all of the targets 14048. For example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record state-related measurements based on vibrations measured with respect to some or all of the targets 14048. In another example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record state-related measurements based on temperatures measured with respect to some or all of the targets 14048. In another example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record state-related measurements based on electrical or magnetic outputs measured with respect to some or all of the targets 14048. In another example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record state-related measurements based on sound outputs measured with respect to some or all of the targets 14048. In another example, some of the mobile

robots and/or the mobile vehicles of the swarm 14038 may be configured to record state-related measurements based on outputs other than vibrations, temperatures, electrical or magnetic, or sound, as measured with respect to some or all of the targets 14048.

[1686] Alternatively, or additionally, different mobile robots and/or mobile vehicles of the swarm 14038 may be configured to record some or all state-related measurements of certain types of the targets 14048. For example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record some or all state-related measurements from agitators (e.g., turbine agitators), airframe control surface vibration devices, catalytic reactors, compressors, and the like. In another example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record some or all state-related measurements from conveyors and lifters, disposal systems, drive trains, fans, irrigation systems, motors, and the like. In another example, some of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record some or all state-related measurements from pipelines, electric powertrains, production platforms, pumps (e.g., water pumps), robotic assembly systems, thermic heating systems, tracks, transmission systems, turbines, and the like. In example embodiments, the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to record some or all state-related measurements of certain types of industrial environments. For example, an industrial environment having targets with states measured using the mobile robots and/or the mobile vehicles of the swarm 14038 may include, but is not limited to, a manufacturing environment, a fossil fuel energy production environment, an aerospace environment, a mining environment, a construction environment, a ship environment, a shipping environment, a submarine environment, a wind energy production environment, a hydroelectric energy production environment, a nuclear energy production environment, an oil drilling environment, an oil pipeline environment, any other suitable energy product environment, any other suitable energy routing or transmission environment, any other suitable industrial environment, a factory, an airplane or other aircraft, a distribution environment, an energy source extraction environment, an offshore exploration site, an underwater exploration site, an assembly line, a warehouse, a power generation environment, a hazardous waste environment, and the like.

[1687] The swarm 14038 includes self-organization systems 14052 for causing the mobile robots or mobile vehicles within the swarm 14038 to self-organize (e.g., during data collection operations within the industrial IoT environment). In example embodiments, a data collection system that includes the swarm 14038 (e.g., the data collection system 12004 or any other suitable data collection system) may include self-organization functionality, which can be performed at or by any of the components of the data collection system. In example embodiments, a mobile robot or mobile vehicle of the swarm 14038 can self-organize without assistance from other components and based on, for example, the data sensed by its associated sensors and other knowledge. In example embodiments, the network 14010 can be accessed for the self-organization without assistance from other components and based on, for example, the data sensed by the mobile robots and/or mobile vehicles, or other knowledge. It should be appreciated that any combination or hybrid-type self-organization system can also be embodied. For example, the data collection

system can perform or enable various methods or systems for data collection having self-organization functionality in an industrial IoT environment. These methods and systems can include analyzing a plurality of sensor inputs, for example, received from or sensed by sensors at the mobile robots and/or at the mobile vehicles of the swarm 14038. The methods and systems can also include sampling the received data and self-organizing at least one of: (i) a storage operation of the data (e.g., with respect to the data pool 14050); (ii) a collection operation of sensors that provide the plurality of sensor inputs, and (iii) a selection operation of the plurality of sensor inputs.

[1688] In example embodiments, the self-organization systems 14052 can be used to collectively organize two or more of the mobile robots and/or the mobile vehicles of the swarm 14038. Alternatively, the self-organization systems 14052 can be used to organize individual mobile robots and/or the mobile vehicles of the swarm 14038. For example, the self-organization systems 14052 can control the traversal of each of the mobile robots and each of the mobile vehicles of the swarm 14038 within different regions, sections, or other divided areas of the industrial IoT environment. In example embodiments, there may be other mobile robots with one or more mobile data collectors integrated therein, other mobile vehicles with one or more mobile data collectors integrated therein, other mobile robots with one or more mobile data collectors coupled thereto, and/or other mobile vehicles with one or more mobile data collectors coupled thereto, which collect data for some or all of the targets 14048, but which are not included in the swarm 14038. Such other mobile robots and/or other mobile vehicles may be controlled individually (e.g., outside of the self-organization systems 14052).

[1689] In example embodiments, the swarm 14038 may include intelligent systems 14054 that process the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038 before transmitting those recorded state-related measurements over the network 14010 or any other suitable communication mechanism. For example, some or all of the mobile robots and/or the mobile vehicles of the swarm 14038 may integrate artificial intelligence processes, machine learning processes, and/or other cognitive processes for analyzing the state-related measurements recorded thereby. In example embodiments, the processing by the intelligent systems 14054 of the mobile robots and/or the mobile vehicles of the swarm 14038 may be or be represented within a pre-processing step of the industrial IoT data collection, monitoring and control system 10. For example, certain types of the mobile robots and/or the mobile vehicles of the swarm 14038 may selectively perform pre-processing of the recorded state-related measurements to identify redundant information, irrelevant information, or insignificant information. In another example, certain types of the mobile robots and/or the mobile vehicles of the swarm 14038 may pre-process the recorded state-related measurements in an automated manner, so as to identify redundant information, irrelevant information, or insignificant information. In another example, the pre-processing may be selectively performed for certain types of state-related measurements recorded by any of the mobile robots and/or the mobile vehicles of the swarm 14038 to pre-process the recorded state-related measurements (e.g., to identify redundant information, irrelevant information, or insignificant information). In another example, the pre-processing may be automated for certain types of state-related measurements recorded by

any of the mobile robots and/or the mobile vehicles of the swarm 14038 to pre-process the recorded state-related measurements (e.g., to identify redundant information, irrelevant information, or insignificant information).

[1690] In example embodiments, the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038 may be made available over the network 14010 (e.g., as described with respect to Figure 307) without the need for external networks. The network 14010 may be a MANET (e.g., the MANET 20 shown in Figure. 2 or any other suitable MANET), the Internet (e.g., the Internet 110 shown in Figure 3 or any other suitable Internet), or any other suitable type of network, or any combination thereof. For example, the network 14010 may be used to receive state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038. The network 14010 may then be used to transmit some or all of those received state-related measurements to other components of the data collection system 102. For example, the network 14010 may be used to transmit some or all of the received state-related measurements to the data pool 14050 (e.g., the data pool 60 shown in Figure 2 or any other suitable data pool) for storage of those received state-related measurements. In another example, the network 14010 may be used to transmit some or all of the received state-related measurements to servers 14056 of the environment for industrial IoT data collection (e.g., the servers 14014 shown in Figure 161, or any other suitable server). The servers 14056 may include one or more hardware or software server aspects. For example, the servers 14056 to which the received state-related measurements are transmitted may include intelligent systems 14058 for processing the received state-related measurements. The intelligent systems 14058 may process the received state-related measurements using artificial intelligence processes, machine learning processes, and/or other cognitive processes to identify information within or otherwise associated with the received state-related measurements. In example embodiments, after processing the received state-related measurements, the servers 14056 to which the received state-related measurements are transmitted may transmit the processed information or data indicative of the processed information to other systems (e.g., for storage or analysis). In example embodiments, the data indicative of the processed information from the servers 14056 may include, for example, output or other results of the artificial intelligence processes, machine learning processes, and/or other cognitive processes.

[1691] In example embodiments, a mobile robot or a mobile vehicle of the swarm 14038 may include a computer vision system or otherwise include computer vision functionality. For example, computer vision functionality of the mobile robot or of the mobile vehicle can include hardware and software configured to identify objects in a multi-axial space using image sensing. In example embodiments, the computer vision functionality within the mobile robot or within the mobile vehicle can include functionality for observing visible states of the targets 14048 during the normal operation of the mobile robot or the mobile vehicle. In example embodiments, data processed by the computer vision functionality of the mobile robot or of the mobile vehicle can be input to the intelligent systems 14054 (e.g., for further processing and learning of the targets 14048 and/or of the environment that includes the targets 14048).

[1692] In example embodiments, some or all of the mobile robots and/or the mobile vehicles of the swarm 14038 may include sensor fusion functionality. For example, the sensor fusion functionality may be embodied as the on-device sensor fusion 80. For example, state-related measurements recorded using multiple analog sensors of one or more of the mobile robots and/or the mobile vehicles of the swarm 14038 (e.g., the multiple analog sensors 82 shown in Figure 4 or any other suitable sensor) may be locally or remotely processed using artificial intelligence processes, machine learning processes, and/or other cognitive processes, which may be embodied within the mobile robots and/or the mobile vehicles of the swarm 14038 themselves, the servers 14056, or both. In example embodiments, the sensor fusion functionality may be embodied by a pre-processing step that is performed prior to the artificial intelligence processes, machine learning processes, and/or other cognitive processes. In example embodiments, the sensor fusion functionality may be performed using a MUX. For example, each of the mobile robots and/or the mobile vehicles of the swarm 14038 may include its own MUX for combining state-related measurements recorded using individual sensors of those multiple sensors. In some such embodiments, the MUX may be internal to the mobile robots and/or the mobile vehicles of the swarm 14038. In some such embodiments, the MUX may be external to the mobile robots and/or the mobile vehicles of the swarm 14038.

[1693] In example embodiments, the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038 may be pulled from the mobile robots and/or mobile vehicles by an upstream device (e.g., a client device or other software or hardware aspect used to review, analyze, or otherwise view the state-related measurements). For example, the mobile robots and/or the mobile vehicles of the swarm 14038 may not actively transmit the state-related measurements that are received (e.g., at the servers 14056, the data pool 14050, or any other suitable hardware or software component that receives the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038). Rather, the transmission of the state-related measurements from the mobile robots and/or the mobile vehicles of the swarm 14038 may be caused by commands received at the mobile robots and/or the mobile vehicles of the swarm 14038 (e.g., from servers 14056 or from other hardware or software of the data collection system 102). For example, a data collector of any of the mobile robots and/or the mobile vehicles of the swarm 14038 may be configured to pull state-related measurements recorded using that mobile robot or mobile vehicle. For example, the mobile robots and/or the mobile vehicles of the swarm 14038 may continuously, periodically, or otherwise at multiple times record state-related measurements within the environment for industrial IoT data collection. The data collector may, at fixed intervals, at random times, or otherwise, transmit one or more commands to some or all of the mobile robots and/or the mobile vehicles of the swarm 14038, for example, to pull some or all of the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038 since the last time state-related measurements were pulled therefrom. In another example, the commands, when processed using individual mobile robots and/or the mobile vehicles of the swarm 14038, cause the recorded state-related measurements or data

representative thereof to be transmitted from the mobile robots and/or the mobile vehicles of the swarm 14038.

[1694] In example embodiments, the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038 may be transmitted from the mobile robots and/or the mobile vehicles of the swarm 14038 responsive to requests for those state-related measurements. For example, the self-organization systems 14052 may, at fixed intervals, at random times, or otherwise, transmit a request for recorded state-related measurements to some or all of the mobile robots and/or the mobile vehicles of the swarm 14038. The processors of some or all of the mobile robots and/or the mobile vehicles of the swarm 14038 to which the request is sent may process the request to determine which state-related measurements to transmit. For example, data indicative of a time of a most recent request for recorded state-related measurements may be accessed by those processors. The processors may then compare that time to a time at which the new request is received from the self-organization systems 14052. The processors may then query a data store for state-related measurements recorded between the two times. The processors may then transmit those state-related measurements in response to the request. In another example, the processors may identify a most recent set of state-related measurements recorded using the corresponding mobile robots and/or the mobile vehicles of the swarm 14038 and transmit those state-related measurements in response to the request. In another example, data collectors within the data collection system 10 may transmit the request directly to the mobile robots and/or the mobile vehicles of the swarm 14038. In yet another example, the mobile robots and/or the mobile vehicles of the swarm 14038 may transmit the request to the self-organization systems 14052. The self-organization systems 14052 may process the request to determine select individual mobile robots and/or the mobile vehicles of the swarm 14038 which were used to record the requested state-related measurements. In example embodiments, the collective processing mind 14020 may then transmit certain state-related measurements in response to the request by, for example, querying a storage for some or all of the state-related measurements recorded using those select individual mobile robots and/or the mobile vehicles of the swarm 14038. Alternatively, the self-organization systems 14052 may process the request to determine which of the state-related measurements recorded by some or all of the mobile robots and/or the mobile vehicles of the swarm 14038 to transmit in response to the request (e.g., based on a time of the request). For example, the self-organization systems 14052 can compare the time of the request to a time of a most recent request for recorded state-related measurements. The self-organization systems 14052 can then retrieve the state-related measurements recorded in between those times and transmit the retrieved state-related measurements in response to the request.

[1695] In example embodiments, the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038 may be pushed to an upstream device (e.g., a client device or other software or hardware aspect used to review, analyze, or otherwise view the state-related measurements). For example, the mobile robots and/or the mobile vehicles of the swarm 14038 may actively transmit the state-related measurements that are received (e.g., at the servers 14056, the data pool 14050, or any other suitable hardware or software component that receives

the state-related measurements recorded using the mobile robots and/or the mobile vehicles of the swarm 14038), without such receiving hardware or software component requesting those state-related measurements or otherwise causing the mobile robot or the mobile vehicle to transmit those state-related measurements based on a command. For example, some or all of the mobile robots and/or the mobile vehicles of the swarm 14038 may transmit state-related measurements on a fixed interval, at random times, immediately upon the recording of those state-related measurements, some amount of time after recording those measurements, upon a determination that a threshold number of state-related measurements have been recorded, or at other suitable times. In some such embodiments, the mobile robots and/or the mobile vehicles of the swarm 14038, either by themselves or using the self-organization systems 14052, may push the recorded state-related measurements in response to detecting a near proximity of a data collection router 14062.

[1696] For example, referring next to Figure 166, upon the detection of the target 14048 by a mobile robot or mobile vehicle 14060 (e.g., one or more of the mobile robot with one or more mobile data collectors integrated therein 14040, the mobile vehicle with one or more mobile data collectors integrated therein 14042, the mobile robot with one or more mobile data collectors coupled thereto 14044, or the mobile vehicle with one or more of the mobile data collectors coupled thereto 14046 of the swarm 14038), the mobile robot or mobile vehicle 14060 records state-related measurements of the target 14048 (e.g., vibrations, temperature, electrical or magnetic output, sound output, or the like). The recorded state-related measurements can be transmitted over the network 14010 (e.g., to the data pool 14050, the servers 14056, or another hardware or software component). Alternatively, the recorded state-related measurements can be transmitted to the data collection router 14062, for example, where the network 14010 is unavailable or where the data collection router 14062 is configured to receive and/or pre-process the recorded state-related measurements from the mobile robot or mobile vehicle 14060. The data collection router 14062 may be one of a number of data collection routers 14062 located throughout the environment for industrial IoT data collection. For example, the data collection router 14062 may be a data collection router 14062 configured to transmit state-related measurements specifically recorded for the target 14048.

[1697] Referring next to Figure 167, various aspects of functionality of intelligent systems 14064 used to process output of the mobile robots and/or the mobile vehicles of the swarm 14038 are disclosed. In example embodiments, the intelligent systems 14064 may include a cognitive learning module 14066, an artificial intelligence module 14068, and a machine learning module 14070. The intelligent systems 14064 may include additional or fewer modules. The intelligent systems 14064 may, for example, be the intelligent systems 14054 or the intelligent systems 14058 shown in Figure 165 or any other suitable intelligent systems. Although shown as separate modules, in example embodiments, there may be overlap between some or all of the cognitive learning module 14066, the artificial intelligence module 14068, and the machine learning module 14070. For example, the artificial intelligence module 14068 may include the machine learning module 14070. In another example, the cognitive learning module 14066 may include the artificial intelligence module 14068 (and, in example embodiments, therefore, the machine learning module

14070). The swarm 14038 may include any number of mobile robots and/or mobile vehicles. For example, as shown, the swarm 14038 includes a first mobile robot or first mobile vehicle 14060A, a second mobile robot or second mobile vehicle 14060B, and an Nth mobile robot or Nth mobile vehicle 14060N, where N is a number greater than two. The intelligent systems 14064 receives the output of the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N. In particular, one or more of the modules 14066, 14068, and 14070 of the intelligent systems 14064 receives data generated by and output from one or more of the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N. The output from the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N may, for example, include state-related measurements recorded using the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N, (e.g., state-related measurements of equipment within an environment for industrial IoT data collection). In example embodiments, the output from the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N may be processed by all three of the modules 14066, 14068, and 14070 of the intelligent systems 14064. In example embodiments, the output from the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N may be processed by only one of the modules 14066, 14068, and 14070 of the intelligent systems 14064. For example, the particular one of the modules 14066, 14068, and 14070 of the intelligent systems 14064 to use to process the output from the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N may be selected based on the mobile robot and/or mobile vehicle used to generate that output, the equipment measured in generating that output, the values of the output, other selection criteria, and the like.

[1698] The knowledge base 14036 (e.g., as described with respect to Figure 164) may be updated based on output from the intelligent systems 14064. The knowledge base 14036 represents a library or other set or collection of knowledge related to the environment of the industrial IoT data collection, including equipment within that environment, tasks performed within that environment, personnel having the skill to perform tasks within that environment, and the like. The intelligent systems 14064 can process the state-related measurements recorded using the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N to facilitate knowledge gathering for expanding the knowledge base 14036. For example, the modules 14066, 14068, and 14070 of the intelligent systems 14064 can process those state-related measurements against existing knowledge within the knowledge base 14036 to update or otherwise modify information within the knowledge base 14036. The intelligent systems 14064 may use intelligence and machine learning capabilities (e.g., of the machine learning module 14070 or as described elsewhere in this disclosure) to process state-related measurements and related information based on detected conditions (e.g., conditions informed by the mobile robots and/or mobile vehicles of the swarm 14038 and/or provided as training data) and/or state information (e.g., state information determined by a machine state recognition system that may determine a state, for example, relating to an operational state, an environmental state, a state within a known process or workflow, a state involving a fault or diagnostic condition, and the like). This may include optimization of input selection and configuration based on learning feedback from the learning feedback system, which may include providing training data (e.g., from a host processing system or from other data collection systems

either directly or from the host processing system) and may include providing feedback metrics (e.g., success metrics calculated within an analytic system of the host processing system). Examples of learning feedback systems, data collection systems, and analytic systems are described elsewhere in this disclosure. Thus, the intelligent systems 14064 can be used to update workflows of tasks assigned and performed within the industrial IoT environment based on output from the mobile robots or mobile vehicles 14060A, 14060B, ... 14060N.

[1699] In example embodiments, the intelligent systems 14064, either within one of the modules 14066, 14068, and 14070 or otherwise, may include other intelligence or machine learning aspects. For example, the intelligent systems 14064 may include one or more of a YOLO neural network, a YOLO CNN, a set of neural networks configured to operate on or from a FPGA, a set of neural networks configured to operate on or from a FPGA and GPU hybrid component, a user configurable series and parallel flow for a hybrid neural network (e.g., configuring series and/or parallel flows between neural networks as outputs which can be communicated between such neural networks), a machine learning system for automatically configuring a topology or workflow for a set of hybrid neural networks (e.g., series, parallel, data flows, etc.) based on a training data set which may or may not use manual configurations (e.g., by a human user), a deep learning system for automatically configuring a topology or workflow for a set of hybrid neural networks (e.g., series, parallel, data flows, etc.) based on a training data set of outcomes from industrial IoT processes (e.g., maintenance, repair, service, prediction of faults, optimization of operation of a machine, system of facility, etc.), or other intelligence or machine learning aspects.

[1700] Thus, in example embodiments, the output of the mobile robots and/or mobile vehicles of the swarm 14038 may be processed using the intelligent systems 14054 to add to, remove from, or otherwise modify the knowledge base 14036. For example, the knowledge base 14036 may reflect information to use to perform one or more tasks within the industrial environment in which the targets are located and in which the mobile robots and/or mobile vehicles of the swarm 14038 are used. The output from the mobile robots and/or mobile vehicles of the swarm 14038 can thus be used to increase knowledge as to the nature of issues that arise with respect to the industrial environment, for example, by describing information about the target from which measurements were recorded, a time and/or date at which the measurements were recorded, pre-existing state or other condition information about the target, information about the time required to resolve an issue with respect to a target, information about how to resolve an issue with respect to a target, information indicating an amount of downtime to the target and to other aspects of the respective industrial environment resulting from resolving the issue, an indication of whether the issue should be resolved now or later (or not at all), and the like. The intelligent systems 14054 may process that output to update existing training data. For example, the existing training data can be used to update the machine learning, artificial intelligence, and/or other cognitive functionality for identifying states of targets based on the output of the mobile robots and/or mobile vehicles of the swarm 14038.

[1701] For example, the knowledge base 14036 may include a series of databases or other tables or graphs arranged hierarchically based on the target or the area of the industrial environment that

includes the target. For example, a first layer of the knowledge base 14036 may refer to the industrial environment (e.g., a power plant, a manufacturing facility, a mining facility, etc.). A second layer of the knowledge base 14036 may refer to zones within the industrial environment (e.g., zone 1, zone 2, etc., or named zones, as the case may be). A third layer of the knowledge base 14036 may refer to targets within those zones (e.g., within a first zone of a power plant including electrical equipment, this could include alternators, circuit breakers, transformers, batteries, exciters, etc., and, within a second zone of a power plant including a turbine, a generator, a generator magnet, etc.). The knowledge base 14036 may be updated based on output of the intelligent systems 14054, by manual user data entry, or both.

[1702] For example, the mobile robots and/or mobile vehicles of the swarm 14038 may be deployed to monitor or otherwise traverse different locations (e.g., zones) within a mining facility used to mine and/or process fuel materials (e.g., coal, natural gas, etc.) and/or non-fuel materials (e.g., stone, sand, gravel, gold, silver, etc.). A mobile robot may be deployed to traverse a first zone in which mineral crushing machinery is operating, and a mobile vehicle may be deployed to traverse a second zone in which underground mining equipment is operating. The mobile robot may measure the operating temperatures of the mineral crushing machinery within the first zone, the temperature of areas of the first zone around the mineral crushing machinery, and the like. The mobile robot may further measure the sound output from the mineral crushing machinery, for example, by recording measurements of the sound output from some or all of the machinery. The mobile robot can detect an overheating issue with respect to one of the mineral crushing machines if it records a temperature measurement which, when processed by the intelligent systems 14054 against the data stored in the knowledge base 14036, indicates that the temperature is at a dangerous level. The mobile robot may be instructed to remain at the location of that machine and record new temperature measurements over some period of time (e.g., at fixed intervals or otherwise) to determine whether the machine is actually operating at a dangerously high temperature. If the intelligent systems 14054 detects that the initial high temperature measurement was not representative of the operating temperature of the machine, the intelligent systems 14054 may either not update the knowledge base 14036 to reflect the misrepresentative measurement or instead may update the knowledge base 14036 to reflect that such a temperature reading may not represent a dangerous condition.

[1703] The mobile vehicle may measure vibrational output with respect to the underground mining equipment. The output of the mobile vehicle may be processed using the intelligent systems 14054 to determine whether it is consistent with the data within the knowledge base 14036 or is deviant therefrom. In the event the output of the mobile vehicle deviates from the data within the knowledge base, the intelligent systems 14054 can update the data within that portion of the knowledge base 14036 to reflect the output of the mobile vehicle. The intelligent systems 14054 may also or instead cause the mobile vehicle to emit an alarm (e.g., using lights, sounds, or both) to warn personnel located in that zone. For example, the intelligent systems 14054 may retrieve information from the knowledge base 14036 suggesting that the output of the mobile vehicle reflects a dangerous condition, for example, related to a potential underground cave-in. In some

scenarios, the intelligent systems 14054 may transmit a notification directly to an operator of the underground machinery to alert them to the dangerous condition.

[1704] A number of handheld devices 14072 are located within the environment for industrial IoT data collection. The handheld devices 14072 may be handheld devices issued by an operator of the environment for industrial IoT data collection. Alternatively, the handheld devices 14072 may be handheld devices owned by workers selected to perform tasks within the environment for industrial IoT data collection. As shown in Figure 168, the handheld devices 14072 include a single handheld device with a single sensor 14074, a single handheld device with multiple sensors 14076, a combination of handheld devices each with a single sensor 14078, and a combination of handheld devices each with one or more sensors 14080. However, in example embodiments, the handheld devices 14072 may include different handheld devices. For example, in example embodiments, the handheld devices 14072 may omit the combination of handheld devices each with the single sensor 14078 and/or the combination of handheld devices each with one or more of the sensors 14080. For example, the handheld devices 14072 may be limited to individual handheld devices rather than combinations of handheld devices that offer combined, improved or otherwise different functionality compared to each of the constituent handheld devices taken individually. In another example, in example embodiments, the handheld devices 14072 may omit the single handheld device with the single sensor 14074 and/or the single handheld device with multiple sensors 14076. For example, the handheld devices 14072 may be limited to combinations of handheld devices rather than individual devices (e.g., where specific combinations of the handheld devices are identified as being valuable in particular contexts or otherwise for recording particular state-related measurements within the environment for industrial IoT data collection).

[1705] In example embodiments, different handheld devices 14072 may be configured to record certain types of state-related measurements of some or all of the targets (e.g., devices or equipment) within the environment for industrial IoT data collection. For example, some of the handheld devices 14072 may be configured to record state-related measurements based on vibrations measured with respect to some or all of the targets. In another example, some of the handheld devices 14072 may be configured to record state-related measurements based on temperatures measured with respect to some or all of the targets. In another example, some of the handheld devices 14072 may be configured to record state-related measurements based on electrical or magnetic outputs measured with respect to some or all of the targets. In another example, some of the handheld devices 14072 may be configured to record state-related measurements based on sound outputs measured with respect to some or all of the targets. In another example, some of the handheld devices 14072 may be configured to record state-related measurements based on outputs other than vibrations, temperatures, electrical or magnetic, or sound, as measured with respect to some or all of the targets.

[1706] Alternatively, or additionally, different handheld devices 14072 may be configured to record some or all state-related measurements of certain types of the targets within the environment for industrial IoT data collection. For example, some of the handheld devices 14072 may be configured to record some or all state-related measurements from agitators (e.g., turbine agitators),

airframe control surface vibration devices, catalytic reactors, compressors, and the like. In another example, some of the handheld devices 14072 may be configured to record some or all state-related measurements from conveyors and lifters, disposal systems, drive trains, fans, irrigation systems, motors, and the like. In another example, some of the handheld devices 14072 may be configured to record some or all state-related measurements from pipelines, electric powertrains, production platforms, pumps (e.g., water pumps), robotic assembly systems, thermic heating systems, tracks, transmission systems, turbines, and the like. In example embodiments, the handheld devices 14072 may be configured to record some or all state-related measurements of certain types of industrial environments. For example, an industrial environment having targets with states measured using the handheld devices 14072 may include, but is not limited to, a manufacturing environment, a fossil fuel energy production environment, an aerospace environment, a mining environment, a construction environment, a ship environment, a shipping environment, a submarine environment, a wind energy production environment, a hydroelectric energy production environment, a nuclear energy production environment, an oil drilling environment, an oil pipeline environment, any other suitable energy product environment, any other suitable energy routing or transmission environment, any other suitable industrial environment, a factory, an airplane or other aircraft, a distribution environment, an energy source extraction environment, an offshore exploration site, an underwater exploration site, an assembly line, a warehouse, a power generation environment, a hazardous waste environment, and the like.

[1707] In example embodiments, the state-related measurements using the handheld devices 14072 may be made available over the network 14010 (e.g., as described with respect to Figure 161) without the need for external networks. The network 14010 may be a MANET (e.g., the MANET 20 shown in Figure. 2 or any other suitable MANET n), the Internet (e.g., the Internet 110 shown in Figure 3 or any other suitable Internet), or any other suitable type of network, or any combination thereof. For example, the network 14010 may be used to receive state-related measurements recorded using the handheld devices 14072. The network 14010 may then be used to transmit some or all of those received state-related measurements to other components of the data collection system 102. For example, the network 14010 may be used to transmit some or all of the received state-related measurements to data pool 14084 (e.g., the data pool 60 shown in Figure 2 or any other suitable data pool) for storage of those received state-related measurements. In another example, the network 14010 may be used to transmit some or all of the received state-related measurements to servers 14086 of the environment for industrial IoT data collection (e.g., the servers 14014 shown in Figure 161, or any other suitable server). The servers 14086 may include one or more hardware or software server aspects. For example, the servers 14086 to which the received state-related measurements are transmitted may include intelligent systems 14088 for processing the received state-related measurements. The intelligent systems 14088 may process the received state-related measurements using artificial intelligence processes, machine learning processes, and/or other cognitive processes to identify information within or otherwise associated with the received state-related measurements. In example embodiments, after processing the received state-related measurements, the servers 14086 to which the received state-related

measurements are transmitted may transmit the processed information or data indicative of the processed information to other systems (e.g., for storage or analysis). The data indicative of the processed information from the servers 14086 may include, for example, output or other results of the artificial intelligence processes, machine learning processes, and/or other cognitive processes.

[1708] In example embodiments, some or all of the handheld devices 14072 may include intelligent systems 14082 for processing the state-related measurements recorded using those handheld devices 14072 before transmitting those recorded state-related measurements (e.g., over the network 14010 or any other suitable communication mechanism). For example, some or all of the handheld devices 14072 may integrate artificial intelligence processes, machine learning processes, and/or other cognitive processes for analyzing the state-related measurements recorded thereby. The processing by the intelligent systems 14082 of the handheld devices 14072 may be or be represented within a pre-processing step of the industrial IoT data collection, monitoring and control system 10. For example, the pre-processing may be selectively performed by certain types of the handheld devices 14072 to pre-process the recorded state-related measurements (e.g., to identify redundant information, irrelevant information, or insignificant information). In another example, the pre-processing may be automated for certain types of the handheld devices 14072 to pre-process the recorded state-related measurements (e.g., to identify redundant information, irrelevant information, or insignificant information). In another example, the pre-processing may be selectively performed for certain types of state-related measurements recorded by any of the handheld devices 14072 to pre-process the recorded state-related measurements (e.g., to identify redundant information, irrelevant information, or insignificant information). In another example, the pre-processing may be automated for certain types of state-related measurements recorded by any of the handheld devices 14072 to pre-process the recorded state-related measurements (e.g., to identify redundant information, irrelevant information, or insignificant information).

[1709] In example embodiments, some or all of the handheld devices 14072 may include sensor fusion functionality. For example, the sensor fusion functionality may be embodied as the on-device sensor fusion 80. For example, state-related measurements recorded using multiple analog sensors of one or more of the handheld devices 14072 (e.g., the multiple analog sensors 82 shown in Figure 4 or any other suitable sensor) may be locally or remotely processed using artificial intelligence processes, machine learning processes, and/or other cognitive processes, which may be embodied within the handheld devices 14072 themselves, the servers 14086, or both. The sensor fusion functionality may be embodied by a pre-processing step that is performed prior to the artificial intelligence processes, machine learning processes, and/or other cognitive processes. In example embodiments, the sensor fusion functionality may be performed using a MUX. For example, each of the single handheld devices with multiple sensors 14076 may include its own MUX for combining state-related measurements recorded using different individual sensors of those multiple sensors. In another example, some or all of the individual handheld devices within the combination of handheld devices each with one or more sensors 14080 may include its own MUX for combining state-related measurements recorded using different individual sensors of

those multiple sensors. In some such embodiments, the MUX may be internal to those handheld devices. In some such embodiments, the MUX may be external to those handheld devices.

[1301] The handheld devices 14072 may be controlled by or otherwise used in connection within the host processing system 112 shown in Figure 6 (or any other suitable host system). The host processing system 112 may be locally accessible over the network 14010. Alternatively, the host processing system 112 may be remote (e.g., as embodied in a cloud computing system), may be accessible using one or more network infrastructure elements (e.g., access points, switches, routers, servers, gateways, bridges, connectors, physical interfaces and the like), and/or use one or more network protocols (e.g., IP-based protocols, TCP/IP, UDP, HTTP, BLUETOOTH™, BLUETOOTH LOW ENERGY™, cellular protocols, LTE, 2G, 3G, 4G, 5G, CDMA, TDSM, packet-based protocols, streaming protocols, file transfer protocols, broadcast protocols, multi-cast protocols, unicast protocols, and the like). In example embodiments, the state-related measurements recorded using the handheld devices 14072 may be processed using a network coding system or method, which may be embodied internally or externally with respect to the host processing system 112. For example, the network coding system can process the measurements recorded using the handheld devices 14072 based on the availability of networks for communicating those recorded state-related measurements, based on the availability of bandwidth and spectrum for communicating those recorded state-related measurements, based on other network characteristics, or based on some combination thereof.

[1710] In example embodiments, the state-related measurements recorded using the handheld devices 14072 may be pulled from the handheld devices 14072 by an upstream device (e.g., a client device or other software or hardware aspect used to review, analyze, or otherwise view the state-related measurements). For example, the handheld devices 14072 may not actively transmit the state-related measurements that are received (e.g., at the servers 14086, the data pool 14084, or any other suitable hardware or software component that receives the state-related measurements recorded using the handheld devices 14072). Rather, the transmission of the state-related measurements from the handheld devices 14072 may be caused by commands received at the handheld devices 14072 (e.g., from servers 14086 or from other hardware or software of the data collection system 102). For example, a data collector, which may be fixed within a particular location of the environment of industrial IoT data collection or mobile therein, may be configured to pull state-related measurements recorded using various handheld devices 14072. For example, the handheld devices 14072 may continuously, periodically, or otherwise at multiple times record state-related measurements within the environment for industrial IoT data collection. The data collector may, at fixed intervals, at random times, or otherwise, transmit one or more commands to some or all of the handheld devices 14072 to pull some or all of the state-related measurements recorded using those handheld devices 14072 since the last time state-related measurements were pulled therefrom. Alternatively, the data collector may, at those fixed intervals, at those random times, or otherwise, transmit the one or more commands to a collective processing mind 14090 associated with the handheld devices 14072. For example, the collective processing mind 14090 may be or include a hub for receiving the state-related measurements recorded using some or all

of the handheld devices 14072. In another example, the commands, when processed using individual handheld devices 14072 or by the collective processing mind 14090 of the handheld devices 14072, cause the recorded state-related measurements or data representative thereof to be transmitted from the handheld devices 14072. For example, the collective processing mind 14090 may be configured to pull the state-related measurements from some or all of the handheld devices 14072 (e.g., at fixed intervals, at random times, or otherwise). The collective processing mind 14090 may then transmit the state-related measurements pulled from the handheld devices 14072 (e.g., to the servers 14086, the data pool 14084, or the other hardware or software component selected or otherwise configured to receive the state-related measurements).

[1711] In example embodiments, the state-related measurements recorded using the handheld devices 14072 may be transmitted from the handheld devices 14072 responsive to requests for those state-related measurements. For example, the collective processing mind 14090 may, at fixed intervals, at random times, or otherwise, transmit a request for recorded state-related measurements to some or all of the handheld devices 14072. The processors of the some or all of the handheld devices 14072 to which the request is sent may process the request to determine which state-related measurements to transmit. For example, data indicative of a time of a most recent request for recorded state-related measurements may be accessed by those processors. The processors may then compare that time to a time at which the new request is received from the collective processing mind 14090. The processors may then query a data store for state-related measurements recorded between the two times. The processors may then transmit those state-related measurements in response to the request. In another example, the processors may identify a most recent set of state-related measurements recorded using the corresponding handheld devices 14072 and transmit those state-related measurements in response to the request. In another example, data collectors within the data collection system 10 may transmit the request directly to the handheld devices 14072. In yet another example, the data collectors may transmit the request to the collective processing mind 14090. The collective processing mind 14090 may process the request to determine select individual handheld devices 14072 which were used to record the requested state-related measurements. The collective processing mind 14090 may then transmit certain state-related measurements in response to the request by, for example, querying a storage for some or all of the state-related measurements recorded using those select individual handheld devices 14072. Alternatively, the collective processing mind 14090 may process the request to determine which of the state-related measurements recorded by some or all of the handheld devices 14072 to transmit in response to the request (e.g., based on a time of the request). For example, the collective processing mind 14090 can compare the time of the request to a time of a most recent request for recorded state-related measurements. The collective processing mind 14090 can then retrieve the state-related measurements recorded in between those times and transmit the retrieved state-related measurements in response to the request.

[1712] In example embodiments, the state-related measurements recorded using the handheld devices 14072 may be pushed from the handheld devices 14072 to an upstream device (e.g., a client device or other software or hardware aspect used to review, analyze, or otherwise view the

state-related measurements). For example, the handheld devices 14072 may actively transmit the state-related measurements that are received (e.g., at the servers 14086, the data pool 14084, or any other suitable hardware or software component that receives the state-related measurements recorded using the handheld devices 14072), without such receiving hardware or software component requesting those state-related measurements or otherwise causing the handheld device to transmit those state-related measurements based on a command. For example, some or all of the handheld devices 14072 may transmit state-related measurements on a fixed interval, at random times, immediately upon the recording of those state-related measurements, some amount of time after recording those measurements, upon a determination that a threshold number of state-related measurements have been recorded, or at other suitable times. In some such embodiments, the handheld devices 14072, either by themselves or using the collective processing mind 14090, may push the recorded state-related measurements in response to detecting a near proximity of a data collection router 14092.

[1713] For example, referring next to Figure 169, the collective processing mind 14090 may include a detector 14094 configured to detect a near proximity of a target 14096 (e.g., one of the devices 13006 shown in Figure 134 or any other suitable target) with respect to one or more of the handheld devices 14072. For example, upon such a detection, the collective processing mind 14090 may send a signal to the one or more of the handheld devices 14072 to record and transmit state-related measurements of receipt at the data collection router 14092. Alternatively, upon such a detection, the collective processing mind 14090 may query a data store to retrieve state-related measurements and then transmit those state-related measurements of receipt at the data collection router 14092. In either case, the data collection router 14092 forwards the received state-related measurements to the servers 14086, the data pool 14084, or any other suitable hardware or software component. In another example, upon such a detection, the collective processing mind 14090 may send the signal directly to the servers 14086, the data pool 14084, or the other hardware or software component, for example, to bypass the data collection router 14092 or where the data collection router 14092 is omitted.

[1714] Referring next to Figure 170, in example embodiments, the collective processing mind 14090 may be omitted. Instead, the handheld devices 14072 detect the near proximity of the target 14096. Upon such detection using the handheld devices 14072 (e.g., one or more of the single handheld device with the single sensor 14074, the single handheld device with multiple sensors 14076, the combination of handheld devices each with the single sensor 14078, or the combination of handheld devices each with one or more sensors 14080), the handheld devices 14072 record state-related measurements of the target 14096 (e.g., vibrations, temperature, electrical or magnetic output, sound output, or the like). The recorded state-related measurements can be transmitted over the network 14010 (e.g., to the data pool 14084, the servers 14086, or any other suitable hardware or software component). Alternatively, the recorded state-related measurements can be transmitted to the data collection router 14092, for example, where the network 14010 is unavailable or where the data collection router 14092 is configured to receive and/or pre-process the recorded state-related measurements from the handheld devices 14072. The data collection router 14092 may be

one of a number of data collection routers 14092 located throughout the environment for industrial IoT data collection. For example, the data collection router 14092 may be a data collection router 14092 configured to transmit state-related measurements specifically recorded for the target 14096.

[1715] Referring next to Figure 171, various aspects of functionality of intelligent systems 14098 used to process output of the handheld devices 14072 are disclosed. The intelligent systems 14098 include a cognitive learning module 14100, an artificial intelligence module 14102, and a machine learning module 14104. In example embodiments, the intelligent systems 14098 may include additional or fewer modules. The intelligent systems 14098 may, for example, be the intelligent systems 14082 or the intelligent systems 14088 shown in Figure 161 or any other suitable intelligent system. Although shown as separate modules, in example embodiments, there may be overlap between some or all of the cognitive learning module 14100, the artificial intelligence module 14102, and the machine learning module 14104. For example, the artificial intelligence module 14102 may include the machine learning module 14104. In another example, the cognitive learning module 14100 may include the artificial intelligence module 14102 (and, in example embodiments, therefore, the machine learning module 14104). The handheld devices 14072 may include any number of handheld devices. For example, as shown, the handheld devices 14072 include a first handheld device 14072A, a second handheld device 14072B, and an Nth handheld device 14072N, where N is a number greater than two. The intelligent systems 14098 receives the output of the handheld devices 14072A, 14072B, ... 14072N. In particular, one or more of the modules 14100, 14102, and 14104 of the intelligent systems 14098 receives data generated by and output from one or more of the handheld devices 14072A, 14072B, ... 14072N. The output from the handheld devices 14072A, 14072B, ... 14072N may, for example, include state-related measurements recorded using the handheld devices 14072A, 14072B, ... 14072N, for example, state-related measurements of equipment within an environment for industrial IoT data collection. In example embodiments, the output from the handheld devices 14072A, 14072B, ... 14072N may be processed by all three of the modules 14100, 14102, and 14104 of the intelligent systems 14098. In example embodiments, the output from the handheld devices 14072A, 14072B, ... 14072N may be processed by only one of the modules 14100, 14102, and 14104 of the intelligent systems 14098. For example, the particular one of the modules 14100, 14102, and 14104 of the intelligent systems 14098 to use to process the output from the handheld devices 14072A, 14072B, ... 14072N may be selected based on the handheld device used to generate that output, the equipment measured in generating that output, the values of the output, other selection criteria, and the like.

[1716] The knowledge base 14036 (e.g., as shown in Figure 164) may be updated based on output from the intelligent systems 14098. The knowledge base 14036 represents a library or other set or collection of knowledge related to the environment of the industrial IoT data collection, including equipment within that environment, tasks performed within that environment, personnel having the skill to perform tasks within that environment, and the like. The intelligent systems 14098 can process the state-related measurements recorded using the handheld devices 14072A, 14072B, ... 14072N to facilitate knowledge gathering for expanding the knowledge base 14036. For example, the modules 14100, 14102, and 14104 of the intelligent systems 14098 can process those state-

related measurements against existing knowledge within the knowledge base 14036 to update or otherwise modify information within the knowledge base 14036. The intelligent systems 14098 may use intelligence and machine learning capabilities (e.g., of the machine learning module 14104 or as described elsewhere in this disclosure) to process state-related measurements and related information based on detected conditions (e.g., conditions informed by the handheld devices 14072 and/or provided as training data) and/or state information (e.g., state information determined by a machine state recognition system that may determine a state, for example, relating to an operational state, an environmental state, a state within a known process or workflow, a state involving a fault or diagnostic condition, and the like). This may include optimization of input selection and configuration based on learning feedback from the learning feedback system, which may include providing training data (e.g., from a host processing system or from other data collection systems either directly or from the host processing system) and may include providing feedback metrics (e.g., success metrics calculated within an analytic system of the host processing system). Examples of host processing systems, learning feedback systems, data collection systems, and analytic systems are described elsewhere in this disclosure. Thus, the intelligent systems 14098 can be used to update workflows of tasks assigned and performed within the industrial IoT environment based on output from the handheld devices 14072A, 14072B, ... 14072N.

[1717] In example embodiments, the intelligent systems 14098, either within one of the modules 14100, 14102, and 14104 or otherwise, may include other intelligence or machine learning aspects. For example, the intelligent systems 14098 may include one or more of a YOLO neural network, a YOLO CNN, a set of neural networks configured to operate on or from a FPGA, a set of neural networks configured to operate on or from a FPGA and GPU hybrid component, a user configurable series and parallel flow for a hybrid neural network (e.g., configuring series and/or parallel flows between neural networks as outputs which can be communicated between such neural networks), a machine learning system for automatically configuring a topology or workflow for a set of hybrid neural networks (e.g., series, parallel, data flows, etc.) based on a training data set which may or may not use manual configurations (e.g., by a human user), a deep learning system for automatically configuring a topology or workflow for a set of hybrid neural networks (e.g., series, parallel, data flows, etc.) based on a training data set of outcomes from industrial IoT processes (e.g., maintenance, repair, service, prediction of faults, optimization of operation of a machine, system of facility, etc.), or other intelligence or machine learning aspects.

[1718] Thus, in example embodiments, the output of the handheld devices 14072 may be processed using the intelligent systems 14088 to add to, remove from, or otherwise modify the knowledge base 14036. For example, the knowledge base 14036 may reflect information to use to perform one or more tasks within the industrial environment in which the targets are located and in which the handheld devices 14072 are used. The output from the handheld devices 14072 can thus be used to increase knowledge as to the nature of issues that arise with respect to the industrial environment, for example, by describing information about the target from which measurements were recorded, a time and/or date at which the measurements were recorded, pre-existing state or other condition information about the target, information about the time required to resolve an

issue with respect to a target, information about how to resolve an issue with respect to a target, information indicating an amount of downtime to the target and to other aspects of the respective industrial environment resulting from resolving the issue, an indication of whether the issue should be resolved now or later (or not at all), and the like. The intelligent systems 14088 may process that output to update existing training data. For example, the existing training data can be used to update the machine learning, artificial intelligence, and/or other cognitive functionality for identifying states of targets based on the output of the handheld devices 14072.

[1719] For example, the knowledge base 14036 may include a series of databases or other tables or graphs arranged hierarchically based on the target or the area of the industrial environment that includes the target. For example, a first layer of the knowledge base 14036 may refer to the industrial environment (e.g., a power plant, a manufacturing facility, a mining facility, etc.). A second layer of the knowledge base 14036 may refer to zones within the industrial environment (e.g., zone 1, zone 2, etc., or named zones, as the case may be). A third layer of the knowledge base 14036 may refer to targets within those zones (e.g., within a first zone of a power plant including electrical equipment, this could include alternators, circuit breakers, transformers, batteries, exciters, etc., and, within a second zone of a power plant including a turbine, a generator, a generator magnet, etc.). The knowledge base 14036 may be updated based on output of the intelligent systems 14088, by manual user data entry, or both. For example, a worker within a manufacturing facility may be given one or more handheld devices (e.g., the handheld devices 14072). The worker may walk around the manufacturing facility and approach several pieces of machinery in different zones, including a hydraulic press within a first zone, a thermoforming machine within a second zone, and a conveyor within a third zone. In approaching the first zone, the handheld device may record a measurement with respect to the hydraulic press indicating a vibration resulting from the operation of the hydraulic press. That measurement is then processed using the intelligent systems 14088, for example, against data stored in a database for the hydraulic press within the knowledge base 14036. In the event the measurement is inconsistent with the data stored in that database, the intelligent system 14088 may determine that the hydraulic press is not operating properly. For example, if the vibration resulting from the operation of the hydraulic press is less than what is recorded in the knowledge base 14036, it may be determined that the hydraulic press is not functioning at an optimal rate. The data within the knowledge base 14036 may then be consulted to determine the likely causes of this issue, including how much time would be required to resolve it. For example, the knowledge base 14036 can indicate that low vibration output is caused by a particular part failure with respect to the hydraulic press.

[1720] The worker may then walk to the thermoforming machine and use the handheld device to measure an ambient temperature around that machine. The measurement is processed using the intelligent systems 14088 to determine that the thermoforming machine is outputting an expected temperature. The worker may then walk to the conveyor and use the handheld machine to measure the velocity of the conveyor. For example, a camera vision system built into the handheld device may be used to detect an operating velocity of the conveyor. The operating velocity may then be compared against the expected operating velocity for the conveyor as shown in the appropriate

section of the knowledge base 14036. Upon a determination that the conveyor is operating at an unexpected velocity, the intelligent systems 14088, such as through the handheld device or through a collective processing mind in communication with the handheld device (e.g., the collective processing mind located within the third zone of the manufacturing facility) may alert workers in the area of the conveyor that the conveyor may not be functioning as intended. The alert may be represented as a warning notification so as to prevent sudden emergency action from being taken. In such a scenario, a worker may see the alert and update the knowledge base 14036 to reflect the unexpected velocity measurement.

[1721] Disclosed herein are systems for using handheld devices for data collection in an industrial environment. As used herein, handheld device integration refers to using handheld devices for specific or general purposes. For example, handheld device integration as described with respect to the functionality or configuration of a system refers to the use by that system of the handheld devices 14072 and/or the hardware and/or software used in connection with the handheld devices 14072 for data collection within an industrial IoT environment, as shown in FIGS. 168 to 171. Such use of handheld devices refers to the use of one or more of the handheld devices 14072. For example, a system disclosed herein as using a handheld device may include using one or more of a mobile phone, laptop computer, tablet computer, personal digital assistant, walkie-talkie, radio, long or short range communication device, flashlight, or other types of handheld devices.

[1722] Systems and methods for identifying operating characteristics, such as vibration, of one or more targets, as described and which may be referred to herein as devices, within an industrial IoT environment using image data sets are described with respect to FIGS. 172 – 174. In example embodiments, a system, such as a computer vision system 15000 generally illustrated in Figure. 172, is configured to detect vibration or other operating characteristics (e.g., vibration, heat, electromagnetic emissions, or other suitable operating characteristics) of the one more targets in the industrial IoT environment (e.g., as described above) using one or more image data sets. The one or more targets may include the devices 13006, as described above. The devices 13006 may include agitators, including turbine agitators, airframe control surface vibration devices, catalytic reactors and compressors. The devices 13006 may also include conveyors and lifters, disposal systems, drive trains, fans, irrigation systems and motors.

[1723] The devices 13006 may also include pipelines, electric powertrains, production platforms, pumps (e.g., water pumps), robotic assembly systems, thermic heating systems, tracks, transmission systems and turbines. The devices 13006 may operate within a single industrial environment 13018 or multiple industrial environments 13018. For example, a pipeline device may operate within an oil and gas environment, while a catalytic reactor may operate in either an oil and gas production environment or a pharmaceutical environment. In example embodiments, an operator, as described throughout this disclosure, operating, supervising, inspecting, or a combination thereof, one or more of the devices 13006 may use the computer vision system 15000 to analyze the operation of the one or more devices 13006. In example embodiments, the operator may review data, reports, charts, or other suitable output from the computer vision system 15000 to determine whether maintenance, repair, or other suitable interaction with the one or more

devices 13006 is required. For example, the output from the computer vision system 15000 may indicate that vibration associated with one of the devices 13006 may lead to a failure if a particular component of the device 13006 is not replaced or repaired within a particular timeframe. In example embodiments, the computer vision system 15000 may be configured to analyze image data sets, as will be described, and identify one or more issues (e.g., faults or potential failures of one or more components), determine a corrective action (e.g., alter an operating speed of a device associated with the faulty or failing component), and initiate the corrective action (e.g., automatically analyze data, identify issues, determine corrective action, and carry out, at least part of, the corrective action).

[1724] A computer vision system, such as the computer vision system 15000, may be adapted to automate tasks and/or features of human visual systems. For example, the computer vision system 15000 may be configured to capture image data associated with the devices 13006 and analyze the image data using various visual techniques that simulate and improve on aspects of human sight and analysis. For example, unlike human sight, the computer vision system 15000 may enhance an image by zooming in on an object, analyzing individual frames and deltas between frames. In another example, the computer vision system 15000 may also capture images outside the typical human perceptible range, such as ultra-violet or infra-red signals. The computer vision system 15000 may then identify various characteristics of the devices 13006, such as the presence or amount of undesirable vibration, using the visual techniques. The computer vision system 15000 may be trained, such as by a human operator or supervisor, or based on a data set, model, or the like. Training may include presenting the computer vision system 15000 with one or more training data sets that represent values, such as sensor data, event data, parameter data, and other types of data (including the many types described throughout this disclosure), as well as one or more indicators of an outcome, such as an outcome of a process, an outcome of a calculation, an outcome of an event, an outcome of an activity, or the like. Training may include training in optimization, such as training the computer vision system 15000 to optimize one or more systems based on one or more optimization approaches, such as Bayesian approaches, parametric Bayes classifier approaches, k-nearest-neighbor classifier approaches, iterative approaches, interpolation approaches, Pareto optimization approaches, algorithmic approaches, and the like. Feedback may be provided in a process of variation and selection, such as with a genetic algorithm that evolves one or more solutions based on feedback through a series of rounds. Feedback may be determined and provided by a human operator or by another component of a monitoring system.

[1725] In example embodiments, the computer vision system 15000 may be trained using training data sets that include visual and/or non-visual data to identify operating characteristics of the devices 13006 using the data captured by one or more data capture devices 15002. In example embodiments, the training data sets may include image data corresponding to various operating states of components of the devices 13006. For example, the training data sets may include image data corresponding to components of the devices 13006 operating within expected or acceptable conditions or tolerances, image data corresponding to components of the devices 13006 operating beyond the expected or acceptable conditions or tolerances, image data corresponding to

components of the devices 13006 operating within the expected or acceptable conditions or tolerances, but are trending toward not operating within the expected or acceptable conditions or tolerances.

[1726] In example embodiments, the training data sets may be generated based on image data of the components of the devices 13006 or similar devices and data captured various sensors (e.g., vibration sensors as described throughout this disclosure). For example, the training data sets may include a correlation of image data with sensed vibrations of components of the devices 13006 (e.g., image data indicating a component is operating within the expected or acceptable conditions or tolerances may be correlated with sensed vibration data that indicates the vibration is expected or acceptable).

[1727] In example embodiments, the computer vision system 15000 may capture data from the devices 13006 (e.g., image data), using various visual input devices. For example, the data capture devices 15002 may capture data, such as visual or image data, during operation of the devices 13006. For example, the data captures devices 15002 may capture a plurality of images over a period of time (e.g., during which the devices 13006 are operating). The data capture devices 15002 may capture images of the devices 13006 at any suitable interval during the period. For example, the data capture devices 15002 may capture an image once per second, once per a fraction of a second, or any suitable interval during the period. In example embodiments, the data capture devices 15002 may capture raw image data. Raw image data may include a signal image, a partial image, data points that represent an image, or other suitable raw image data. In example embodiments, the data capture devices 15002 may encode the raw image data using any suitable image encode techniques.

[1728] The data capture devices 15002 may include cameras, sensors, other image capture devices, other data capture devices, or a combination thereof. In example embodiments, the data capture devices 15002 may include one or more full spectrum cameras configured to capture image data that includes visible light image data and/or non-visible light image data, including infrared image data, ultraviolet image data, other non-visible image data, or a combination thereof. In example embodiments, the data capture devices 15002 may include one or more radiation imaging devices, such as an X-ray imaging device or other suitable radiation imaging device. The one or more radiation imaging devices may be configured to capture image data of the devices 13006 during operation of the devices 13006 using X-ray imaging or other suitable radiation imaging. In example embodiments, the data capture devices 15002 may include one or more sonic capture device configured to capture image data of the devices 13006 during operation of the devices 13006 using sound waves, such as ultrasonic sound waves or other suitable sound waves. In example embodiments, the data capture devices 15002 may include a light imaging, detection, and ranging (LIDAR) device configured to capture image data of the devices 13006 during operation of the devices 13006 by measuring the distance to a target by illuminating the target with a pulsed light and measuring the reflected pulses with one or more sensors. In example embodiments, the data capture devices 15002 may include a point cloud data capture device configured to capture image

data of the devices 13006 during operation of the devices 13006 using lasers or other suitable light to generate a set of data points represent a 3-dimensional model of the devices 13006.

[1729] In example embodiments, the data capture devices 15002 may include an infrared inspection device configured to capture image data of the devices 13006 during operation of the devices 13006 using infrared imaging. In example embodiments, the data capture devices 15002 may include a digital image capturing device, such as a digital camera, configured to capture image data of the devices 13006 during operation of the devices 13006 using visible light. For example, an operator operating, supervising, monitoring, and/or inspecting one or more of the devices 13006 may utilize a mobile device, such as a mobile phone, smart phone, tablet computer, or other suitable mobile device. The mobile device may include an image capture device, such as a digital camera. The operator may capture image data associated with the image capture device of the mobile device. In example embodiments, the data capture device 15002 may be a stand-alone device that captures image data, as described, and communicates the captured image data to a client, a server, or a combination thereof, as will be described.

[1730] In example embodiments, one or more data capture devices 15002 may be positioned at or near a respective device 13006 at predefined distances and locations with respect to the respective device 13006. The predefined distances and locations at which the one or more data capture devices 15002 are positioned, or disposed, may be selected such that the one or more of the data capture devices 15002 has a desired field of data capture of a point of interest of the respective device 13006. The point of interest may include any suitable point or areas of the respective device 13006. For example, the point of interest may include a belt, bearing, blade, vane, fan, or any other suitable component, point or area of interest on or related to the respective device 13006. The field of data capture may include a field of vision for an image data capture device 15002, a field of sonic data capture for a sonic data capture device 15002, or other suitable field of data capture. The data captured from the combine fields of data capture from each respective data capture device positioned at or near the respective device 13006 may be used, as will be described, by the image data set generator 15006 to generate one or more image data sets that represent images of the point of interest of the respective device 13006. In example embodiments, the data capture devices 15002 may include any combination of the devices described herein or other suitable data capture devices not described.

[1731] In example embodiments, the data capture devices 15002 may capture image data of the devices 13006, as described, and communicate the captured image data to a client 15004 and/or a server 15010 using a network 15008. The client 15004 may include any suitable client including those described throughout this disclosure. In example embodiments, the client 15004 may be a mobile device, or other suitable clients. The client may include a processor configured to execute instructions (e.g., instructions that, when executed by the processor, cause the processor to execute various portions of the computer vision system 15000 or various methods described herein) stored on a memory. The client 15004 may be owned, operated, and/or utilized by an operator working on or near the devices 13006, as described throughout this disclosure. The network 15008 may be any suitable network, including any network described throughout this disclosure, including, but

not limited to, the Internet, a cloud network, a local area network, a wide area network, a wireless network, a wired network, a cellular network, and the like, or any combination thereof. The server 15010 may be any suitable server, including any server described throughout this disclosure. The server 15010 may include a processor configured to execute instructions (e.g., instructions that, when executed by the processor, cause the processor to execute various portions of the computer vision system 15000 or various methods described herein) stored on a memory. The server 15010 may be a stand-alone server or a group of servers. The server 15010 may be a dedicated server or one of a distributed computing servers or a cloud server, and the like, or any combination thereof.

[1732] In example embodiments, the computer vision system 15000 may include an image data set generator 15006. The image data set generator 15006 may comprise an application or other suitable software or program capable of being executed on the client 15004 and/or the server 15010. In example embodiments, the client 15004 may be configured to execute the image data set generator 15006. For example, an operator, as described, may carry the client 15004 as the operator interacts with a first device 13006. One or more of the data capture devices 15002 may be configured to capture image data, as described, associated with the first device 13006. For example, a first data capture device 15002 may be disposed near the first device 13006, such that, the first data capture device 15002 has a field of data capture, as described, to a point of interest on the first device 13006. The first data capture device 15002 may capture raw image data associated with the first device 13006. The first data capture device 15002 may communicate, via the network 15008, the raw image data to the client 15004. The image data set generator 15006 may generate one or more image data sets, as will be described, using the raw image data. In example embodiments, the server 15010 may be configured to execute the image data set generator 15006, as is generally illustrated in Figure 152. The first data capture device 15002 may communicate, via the network 15008, the raw image data to the server 15010. The image data set generator 15006, being executed by the server 15010, may generate one or more image data sets, as will be described, using the raw image data.

[1733] In example embodiments, the image data set generator 15006 may be configured to generate one or more image data sets using raw image data received from the one or more data capture devices 15002. The image data sets may include images that include data capable (e.g., in a suitable format) of being analyzed or processed by the vision analytics module 15012, as will be described. The image data set generator 15006 may be configured to decode raw image data. For example, as described, the one or more data capture devices 15002 may encode raw image data before communicating the encoded raw image data to the client 15004 and/or the server 15010. The image data set generator 15006 may be configured to decode the raw image data using any suitable image decoding techniques. In example embodiments, the image data set generator 15006 may be configured to correlate related raw image data, stitch raw image data (e.g., by using multiple images from one or more data capture devices 15002 to create a single image of a point of interest on one of the devices 13006), or generate image data sets using any suitable image data set generation techniques, and/or any suitable image processing techniques.

[1734] In example embodiments, the image data set generator 15006 may generate the image data sets from raw data comprising data other than visible light image data. For example, as described, the data capture devices 15002 may capture data such as sonic data, non-visible light data, and other various data. The image data set generator 15006 may receive the non-image raw data and convert the non-image raw data into image data. For example, the image data set generator 15006 may generate one or more images of the point of interest of the device 13006 using sound waves captured by one or more data capture devices 15002. The image data set generator 15006 may generate the image data set using any suitable technique. The image data set generator 15006 may communicate the one or more image data sets to a vision analytics module 15012.

[1735] In example embodiments, the vision analytics module 15012 may be an application or other suitable software capable of being executed on the server 15010. While the vision analytics module 15012 is illustrated and described as being executed by the server 15010, it should be understood that the client 15004 may be configured to execute the vision analytics module 15012.

[1736] As is generally illustrated in Figure 174, the vision analytics module 15012 may include an image data database 15014, a training data database 15016, a visual analyzer 15018, and an operating characteristics detector 15020. In example embodiments, the image data database 15014 may include any suitable database and may be disposed locally on the client 15004 and/or the server 15010, remotely from either of the client 15004 and the server 15010, or other suitable location. The image data database 15014 may store the image data sets generated by the image data set generator 15006, as described. For example, the image data set generator 15006 may generate one or more image data sets, as described, and communicate the one or more image data sets to the image data database 15014. In example embodiments, the image data database 15014 may be any suitable image repository configured to store the image data sets.

[1737] The training data database 15016 may include any suitable database and may be disposed locally on the client 15004 and/or the server 15010, remotely from either of the client 15004 and the server 15010, or other suitable location. The training data database 15016 may store the training data sets generated by a deep learning system, as will be described. In example embodiments, the training data database 15016 may be any suitable training data repository configured to store the training data sets. The training data sets may include any suitable training data sets. For example, the training data sets may be generated by a deep learning system, as will be described, using various suitable image data sets, such as image data sets representing portions of the devices 13006, portions of other devices, image data sets representing motion, vibration, or other various characteristics of the devices 13006 or other devices, or any other suitable image data sets or other data sets.

[1738] In example embodiments, the training data sets may be used to train the computer vision system 15000 to detect the various operating characteristics of the devices 13006. For example, as will be described, the deep learning system may train the visual analyzer 15018 to identify various data points of the image data sets, such as, anomalies, features, characteristics, or other suitable data points. In example embodiments, the visual analyzer 15018 may be trained by any suitable training system, such as a machine learning system, an artificial intelligence training system, deep

learning system, programed by a human programmer, or configured, trained, programed, etc. using any suitable techniques, methods, and/or systems. For example, the visual analyzer 15018 may be configured to identify a portion of a point of interest of a respective device 13006 represented in an image data set. For example, the visual analyzer 15018 may identify a portion of a belt of the respective device 13006 represented by the image data set. The visual analyzer 15018 may be configured to analyze the portion of the point of interest and determine whether the characteristics (e.g., position, size, shape, and/or other suitable characteristics) of the portion of the point of interest corresponds to predicted or predetermined characteristics of the portion of the point of interest. For example, the visual analyzer 15018 may identify the portion of the point of interest in one of a plurality of images associated with the image data set. The visual analyzer 15018 may record values corresponding to various characteristics of the portion of the point of interest associated with each of the plurality of images of the image data set. For example, the visual analyzer 15018 may record a position of a portion of a belt of the respective device 13006 in each image of the plurality of successive images of the image data set and may track the delta in the position of the belt in the successive images.

[1739] The predicted or predetermined characteristics may be predicted or predetermined based on the training data sets and may correspond to characteristics of the portion for the point of interest where the portion of the point of interest indicates that the respective device 13006 is operating within acceptable or expected tolerances. For example, the predicted or predetermined characteristics of the portion of the point of interest may include a position of a portion of a belt while the respective device 13006 is operating. The position of the belt may correspond to an expected operating position of the belt while the respective device 13006 is operating (e.g., where the portion of the belt is expected to be while the respective device 13006 is operating according to acceptable operating tolerances). While various examples are described, it should be understood that the visual analyzer 15018 may use any suitable characteristics of the portion of the point of interest to analyze the image data sets.

[1740] In example embodiments, the visual analyzer 15018 may compare the recorded characteristics of the portion of the point of interest with the predicted or predetermined characteristics of the portion of the point of interest. The visual analyzer 15018 may be configured (e.g., trained, configured, programmed, etc., as described above), to generate analytics of the portion of the point of interest based on the comparison of the recorded characteristics of the portion of the point of interest with the predicted or predetermined characteristics of the portion of the point of interest. For example, the visual analyzer 15018 may determine a variance between a recorded position of the portion of the point of interest and a predicted or predetermined position of the portion of the point of interest (e.g., a variance between an actual or observed position of, for example, the belt of the respective device 13006 a predicted or predetermined position of the belt of the respective device 13006). As described, the image data set may include a plurality of images of the portion of the point of interest captured over a period. The visual analyzer 15018 may determine a first variance between a first recorded characteristic of the portion of the point of interest and a first predicted or predetermined characteristic of the portion of the point of interest

at a first interval during the period (e.g., using a first image captured during the first interval). The visual analyzer 15018 may then determine a second variance between a second recorded characteristic of the portion of the point of interest and a second predicted or predetermined characteristic of the portion of the point of interest at a second interval during the period (e.g., using a second image captured during the second interval). The visual analyzer 15018 may continue to determine variances for a plurality of recorded characteristics and a plurality of predicted or predetermined characteristics over the period using images corresponding to intervals during the period. In this manner, the visual analyzer 15018 may generate data that represents the variance of the characteristics of the portion of the point of interest with respect to the predicted or predetermined characteristics of the portion of the point of interest overtime. For example, the visual analyzer 15018 may generate data that represents the difference in the actual or observed position of the belt compared to the predicted or predetermined position of the belt over a period of time. The visual analyzer 15018 may quantize the variance. For example, the visual analyzer 15018 may be configured to determine a value representing the variance between the recorded characteristics and the predicted or predetermined characteristics (e.g., a value representing a distance between a recorded position of the belt and a predicted or predetermined position of the belt). In example embodiments, the visual analyzer 15018 may be configured to generate a variance data set that includes values representing the variances between the recorded characteristics of the portion of the point of interest and the predicted or predetermined portion of the point of interest. The visual analyzer 15018 may communicate the variance data set to the operating characteristics detector 15020.

[1741] In example embodiments, the operating characteristics detector 15020 may be located or disposed on the vision analytics module 15012 or located or disposed remotely from the vision analytics module 15012. In example embodiments, the operating characteristics detector 15020 may be configured to determine or identify various operating characteristics of the respective device 13006, or any suitable device 13006, based on the variance data set. The various operating characteristics may include vibration, heat, distortion, deflection, other suitable operating characteristics, or a combination thereof of the portion of the point of interest during operating of the respective device 13006, vibration, heat, distortion, deflection, other suitable operating characteristics, or a combination thereof of other portions of the respective device 13006, other suitable operating characteristics of the respective device 13006, or a combination thereof. As described, the operating characteristics detector 15020 may be trained by any suitable training system, such as a machine learning system, an artificial intelligence training system, deep learning system, programed by a human programmer, or configured, trained, programed, etc. using any suitable techniques, methods, and/or systems. In example embodiments, the operating characteristics detector 15020 may be configured to identify operating characteristics of the portion of the point of interest by identifying various data of the variance data set that indicate quantities or other suitable measurements of one or more operating characteristics of the respective device 13006.

[1742] For example, the operating characteristics detector 15020 may identify data of the variance data set that indicates that the belt is vibrating at a first frequency (e.g., by identifying values associated with the variance data set that indicate that the position of the belt over a period of time is moving at a first frequency). The operating characteristics detector 15020 may compare the identified operating characteristics with trained or programmed operating characteristics to determine whether the operating characteristics are within operating tolerance for the respective device 13006. For example, the operating characteristics detector 15020 may compare a value associated with the operating characteristic with a threshold value (e.g., and determine whether the operating characteristic is within tolerances depending on whether the operating characteristic value is above or below the threshold), compare the value associated with the operating characteristic to a predicted value (e.g., and determine if the values are different that the operating characteristic is not operating within tolerances), or other suitable determinative analysis, or a combination thereof. For example, the operating characteristics detector 15020 may compare the frequency at which the belt is vibrating with a trained or programmed frequency. The trained or programmed frequency may include a frequency of vibration of the belt during normal or acceptable operation of the respective device 13006, a frequency of vibration of the belt that indicates the belt is vibrating beyond acceptable tolerances, a frequency of vibration that is within the normal or acceptable operation of the respective device 13006 and indicates that the belt may eventually vibrate at a frequency beyond the acceptable tolerances of the operation of the respective device 13006, or other suitable frequencies. While only vibration is described, the trained or programed operating characteristics may indicate any suitable operating characteristics of the respective device 13006. The operating characteristics detector 15020 may output (e.g., to a database, to a report, to monitor, or other suitable output location or device) an operatic characteristics data set that includes data indicating values or the operating characteristics and/or information indicating predictive (e.g., future) operating characteristics (e.g., determined based on the actual or observed operating characteristics of the portion of the point of interest and the trained or programed operating characteristic that indicate that the actual or observed operating characteristics indicate particular further operating characteristics), actual or observed operating characteristics, other suitable information or values, or a combination thereof.

[1743] In example embodiments, an operator may review and/or analyze the operating characteristics data set to determine whether the respective device 13006, and/or the portion of the point of interest of the respective device 13006, is operating within expected or acceptable tolerances. Additionally, or alternatively, the operator may determine, based on the operating characteristics data set that one or more components of the respective device 13006 is faulty, will become faulty, requires maintenance, or other suitable determinations. For example, the operating characteristics data set may indicate that the belt is vibrating at a first frequency. The belt vibrating at the first frequency may indicate that a pulley associated with the belt is faulty or requires maintenance. The operator may maintain or replace the pulley based on the operating characteristics data. In example embodiments, the operating characteristics detector 15020 may be configured to output information or data that indicates that a component of the respective device

13006 requires maintenance or replacement. For example, as described, the operating characteristics data set may indicate that the belt is vibrating at the first frequency. The operating characteristics detector 15020 may be configured to determine, based on the operating characteristics data set (e.g., indicating that the belt is vibrating at the first frequency), and the trained or programmed operating characteristics that the belt vibrating at the first frequency indicates that a first pulley is faulty and should be replaced or maintained. The operating characteristics detector 15020 may output the information or data to the operator, as described, who may then act on the information or data (e.g., by replacing or maintaining the first pulley).

[1744] In example embodiments, the computer vision system 15000 may capture data from the respective devices 13006 (e.g., non-image data), using various non-visual input devices. For example, the data capture devices 15002 may capture data, such as temperature, pressure, chemical structure, other suitable non-visual data, or a combination thereof, during operation of the respective devices 13006. A chemical structure may include a molecular geometry representing spatial arrangements of atoms in a molecular and the chemical bonds that hold the atoms together. A chemical structure can be represented by molecular models or formulas. For example, the data captures devices 15002 may capture a plurality of measurement values over a period of time (e.g., during which the respective devices 13006 are operating). The data capture devices 15002 may capture measurements of the respective devices 13006 at any suitable interval during the period. For example, the data capture devices 15002 may capture a measurement once per second, once per a fraction of a second, or any suitable interval during the period. In example embodiments, the data capture devices 15002 may capture raw measurement data. Raw measurement data may include a temperature measurement, a pressure measurement (e.g., of liquid or gas within a portion of the respective device 13006), a chemical structure measurement (e.g., of a liquid, gas, or solid within a portion of the respective device 13006), or other suitable raw measurement data. In example embodiments, the data capture devices 15002 may encode the raw measurement data using any suitable measurement encoding techniques.

[1745] The data capture devices 15002 may include pressure sensors, temperature sensors, chemical sensors, fluid sensors, other sensors, other data capture devices, or a combination thereof. In example embodiments, the data capture devices 15002 may include one or more pressure sensors configured to capture pressure measurement data that includes of a portion of the respective device 13006. For example, a pressure sensor may measure pressure within a vat, pipe, tank, or other suitable pressurized enclosure of the respective device 13006. In example embodiments, the data capture devices 15002 may include one or more temperature sensors configured to measure temperature of a portion of the respective device 13006. For example, a temperature sensor may measure the temperature of oven, kiln, vat, pipe, tank, or other suitable portions of the respective device 13006. In example embodiments, the data capture devices 15002 may include one or more chemical sensors configured to measure or determine a chemical structure of a liquid, gas, or solid associated with the respective device 13006. For example, a chemical sensor may measure the chemical structure of a part manufactured by the respective device 13006, the chemical structure of cooling fluid used to cool the respective device 13006 during operation, the chemical structure

of waste produced by the respective device 13006 during operation, or other suitable chemical structures of other suitable liquids, fluids, gases, or solids associated with the respective device 13006.

[1746] In example embodiments, the data capture devices 15002 may be associated with a mobile device. For example, an operator operating, supervising, monitoring, and/or inspecting one or more of the respective devices 13006 may utilize a mobile device, such as a mobile phone, smart phone, tablet computer, or other suitable mobile device. The mobile device may include a data capture device, such as an add-on sensor. The operator may capture measurement data using the add-on sensor of the mobile device. In example embodiments, the data capture device 15002 may be a stand-alone device that captures measurement data, as described, and communicates the captured measurement data to the client 15004, the server 15010, or a combination thereof, as described.

[1747] In example embodiments, one or more data capture devices 15002 may be positioned at or near a respective device 13006 at predefined distances and locations with respect to the respective device 13006. The predefined distances and locations at which the one or more data capture devices 15002 are positioned, or disposed, may be selected such that the one or more data capture devices 15002 has a desired field of data capture of a point of interest of the respective device 13006. As described, the point of interest may include any suitable point or areas of the respective device 13006. For example, the point of interest may include a vat, tank, pipe, enclosure, manufactured part, coolant fluid, waste product, other suitable points of interest, or a combination thereof. The field of data capture may include an area in which the desired measurement can be captured using the data capture devices 15002. The data captured from the combine fields of data capture from each respective data capture device 15002 positioned at or near the respective device 13006 may be used, as described, by the image data set generator 15006 to generate one or more image data sets that represent images of the point of interest of the respective device 13006. In example embodiments, the data capture devices 15002 may include any combination of the devices described herein or other suitable data capture devices not described.

[1748] In example embodiments, the data capture devices 15002 may capture measurement data of the respective devices 13006, as described, and communicate the captured measurement data to the client 15004 and/or the server 15010 using the network 15008. The client 15004 may include any suitable client including those described throughout this disclosure. In example embodiments, the client 15004 may be a mobile device, or other suitable clients. The client 15004 may be owned, operated, and/or utilized by an operator working on or near the respective devices 13006, as described throughout this disclosure. The network 15008 may be any suitable network, including any network described throughout this disclosure, including, but not limited to, the Internet, a cloud network, a local area network, a wide area network, a wireless network, a wired network, a cellular network, and the like, or any combination thereof. The server 15010 may be any suitable server, including any server described throughout this disclosure. The server 15010 may be a stand-alone server or a group of servers. The server 15010 may be a dedicated server or one of a distributed computing servers or a cloud server, and the like, or any combination thereof.

[1749] In example embodiments, as described, the image data set generator 15006 may comprise an application or other suitable software or program capable of being executed on the client 15004 and/or the server 15010. In example embodiments, the client 15004 may be configured to execute the image data set generator 15006. For example, an operator, as described, may carry the client 15004 as the operator interacts with a first device 13006. One or more of the data capture devices 15002 may be configured to capture measurement data, as described, associated with the first device 13006. For example, a first data capture device 15002 may be disposed near the first device 13006, such that, the first data capture device 15002 has a field of data capture, as described, to a point of interest on the first device 13006. The first data capture device 15002 may capture raw measurement data associated with the first device 13006. The first data capture device 15002 may communicate, via the network 15008, the raw measurement data to the client 15004. The image data set generator 15006 may generate one or more image data sets using the raw measurement data. In example embodiments, the server 15010 may be configured to execute the image data set generator 15006, as is generally illustrated in Figure 152. The first data capture device 15002 may communicate, via the network 15008, the raw measurement data to the server 15010. The image data set generator 15006, being executed by the server 15010, may generate one or more image data sets using the raw measurement data.

[1750] In example embodiments, the image data set generator 15006 may be configured to generate one or more image data sets using raw measurement data received from the one or more data capture devices 15002. The image data sets may include images that include data capable (e.g., in a suitable format) of being analyzed or processed by the vision analytics module 15012, as described. The image data set generator 15006 may be configured to decode raw measurement data. For example, as described, the one or more data capture devices 15002 may encode raw measurement data before communicating the encoded raw measurement data to the client 15004 and/or the server 15010. The image data set generator 15006 may be configured to decode the raw measurement data using any suitable measurement decoding techniques. For example, the image data set generator 15006 may be configured to interpret a signal representing a measured value as the measurement value. In example embodiments, the image data set generator 15006 may be configured to correlate related raw measurement data, stitch raw measurement data (e.g., by using multiple measurements from one or more data capture devices 15002 to create a single value that represents a point of interest on one of the respective devices 13006), or generate image data sets using any suitable image data set generation techniques, and/or any suitable measurement data processing techniques. For example, the image data set generator 15006 may be configured to use measurement data corresponding to pressure, temperature, chemical structure, or other suitable measurement data, to generate image data that represents the point of interest of the respective device 13006.

[1751] In example embodiments, the image data set generator 15006 may be configured to use measurement data, as described, in combination with raw image data (e.g., captured by the data capture devices 15002, as described above), to generate one more image data sets. For example, the image data set generator 15006 may be configured to generate an image of the point of interest

of the respective device 13006 using captured image data combined with an associated temperature measurement to generate a precise image of the point of interest (e.g., accounting for, for example, component expansion, deflection, growth, shrinkage, or other change in shape or size due to the temperature of the component). The image data set generator 15006 may communicate the one or more image data sets to a vision analytics module 15012. In example embodiments, the vision analytics module 15012 may be an application or other suitable software capable of being executed on the server 15010. While the vision analytics module 15012 is illustrated and described as being executed by the server 15010, it should be understood that the client 15004 may be configured to execute the vision analytics module 15012. In example embodiments, the vision analytics module 15012 may analyze the image data sets, as described. For example, the visual analyzer 15018 may analyze the image data sets. The operating characteristics detector 15020 may identify operating characteristics, as described.

[1752] In example embodiments, as described, the training data database 15016 may include any suitable database and may be disposed locally on the client 15004 and/or the server 15010, remotely from either of the client 15004 and the server 15010, or other suitable location. The training data database 15016 may store the training data sets generated by a deep learning system, as will be described. In example embodiments, the training data database 15016 may be any suitable training data repository configured to store the training data sets. The training data sets may include any suitable training data sets. For example, the training data sets may be generated by a deep learning system, as will be described, using various suitable data sets, such as data sets representing portions of the respective devices 13006, portions of other devices, data sets representing pressure, data sets representing temperature, data sets representing chemical structure, data sets representing vibration, or other various characteristics of the respective devices 13006 or other devices, or any other suitable data sets.

[1753] In example embodiments, the training data sets may be used to train the computer vision system 15000 to detect the various operating characteristics of the respective devices 13006. For example, as will be described, the deep learning system may train the visual analyzer 15018 to identify various data points of the image data sets, such as, anomalies, features, characteristics, or other suitable data points. In example embodiments, the visual analyzer 15018 may be trained by any suitable training system, such as a machine learning system, an artificial intelligence training system, deep learning system, programed by a human programmer, or configured, trained, programed, etc. using any suitable techniques, methods, and/or systems. For example, the visual analyzer 15018 may be configured to identify a portion of a point of interest of the respective device 13006 represented in an image data set. For example, the visual analyzer 15018 may identify a portion of a belt of the respective device 13006 represented by the image data set. The visual analyzer 15018 may be configured to analyze the portion of the point of interest and determine whether the characteristics (e.g., position, size, shape, and/or other suitable characteristics) of the portion of the point of interest corresponds to predicted or predetermined characteristics of the portion of the point of interest. For example, the visual analyzer 15018 may identify the portion of the point of interest in one of a plurality of images associated with the image data set. The visual

analyzer 15018 may record various characteristics of the portion of the point of interest associated with each of the plurality of images of the image data set. For example, the visual analyzer 15018 may record a pressure value, a temperature value, or other suitable measured value associated with a portion of a belt of the respective device 13006 in each image of the plurality of successive images of the image data set and may track the delta in the measured values of the belt in the successive images (e.g., using the measured values captured by the data capture devices 15002, as described). As described, the visual analyzer 15018 may generate variance data sets based on the deltas between the recorded values and the predicted or predetermined values.

[1754] In example embodiments, the operating characteristics detector 15020 may be located or disposed on the vision analytics module 15012 or located or disposed remotely from the vision analytics module 15012. In example embodiments, the operating characteristics detector 15020 may be configured to determine or identify various operating characteristics of the respective device 13006, or any suitable respective device 13006, based on the variance data set. The various operating characteristics may include vibration, heat, distortion, deflection, other suitable operating characteristics, or a combination thereof of the portion of the point of interest during operating of the respective device 13006, vibration, heat, distortion, deflection, other suitable operating characteristics, or a combination thereof of other portions of the respective device 13006, other suitable operating characteristics of the respective device 13006, or a combination thereof.

[1755] As described, the operating characteristics detector 15020 may be trained by any suitable training system, such as a machine learning system, an artificial intelligence training system, deep learning system, programed by a human programmer, or configured, trained, programed, etc. using any suitable techniques, methods, and/or systems. In example embodiments, the operating characteristics detector 15020 may be trained by a deep learning system, as will be described, using the training data sets that include data sets representing portions of the respective devices 13006, portions of other devices, data sets representing pressure, data sets representing temperature, data sets representing chemical structure, data sets representing vibration, or other various characteristics of the respective devices 13006 or other devices, or any other suitable data sets. In example embodiments, the operating characteristics detector 15020 may be configured to identify operating characteristics of the portion of the point of interest by identifying various data of the variance data set that indicate quantities or other suitable measurements of one or more operating characteristics of the respective device 13006. In example embodiments, the operating characteristics may include a pressure within a component of the respective device 13006, a temperature of at least a portion of a component of the respective device 13006, a chemical structure of a material (e.g., gas, liquid, or solid of or within a component of the respective device 13006 or of a component or part manufactured by the respective device 13006), a density of a material (e.g., gas, liquid, or solid of or within a component of the respective device 13006 or of a component or part manufactured by the respective device 13006), other suitable operating characteristics, or a combination thereof.

[1756] For example, the operating characteristics detector 15020 may identify data of the variance data set that indicates that a component of the respective device 13006 is misshapen due to an

unexpected increase in temperature (e.g., by identifying values associated with the variance data set that indicate that the temperature of the component over a period of time is increasing at a rate greater than expected). The operating characteristics detector 15020 may compare the identified operating characteristics with trained or programmed operating characteristics to determine whether the operating characteristics are within operating tolerance for the respective device 13006. For example, the operating characteristics detector 15020 may compare the rate of temperature change of the component with a trained or programmed rate of temperature change of the component. The operating characteristics detector 15020 may output (e.g., to a database, to a report, to monitor, or other suitable output location or device) an operatic characteristics data set that includes data indicating values or the operating characteristics and/or information indicating predictive (e.g., future) operating characteristics (e.g., determined based on the actual or observed operating characteristics of the portion of the point of interest and the trained or programed operating characteristic that indicate that the actual or observed operating characteristics indicate particular further operating characteristics), actual or observed operating characteristics, other suitable information or values, or a combination thereof. As described, an operator may analyze the output data and take appropriate corrective action. Additionally, or alternatively, the computer vision system 15000 may automatically identify a corrective action and initiate the corrective action.

[1757] In example embodiments, the computer vision system 15000 may implement a classification model (e.g., using a deep neural network, or other suitable neural or other networks). For example, the vision analytics module 15012 may implement a classification module that receives analytics of the image data, including the variance data sets described above. The vision analytics module 15012 may output a classification relating to an operating characteristic of the respective device 13006. For example, the classification model, via the vision analytics module 15012, may receive features defining the variances between the recorded characteristics of the image data sets of the belt of the respective device 13006, in operation. The classification model, having been trained using image data and/or non-image data corresponding to faulty belts, image data and/or non-image data corresponding to belts not yet faulty, and image and/or non-image data corresponding to belts operating in an expected and/or acceptable condition, may output a classification that indicates whether the belt is faulty, operating within expected and/or acceptable condition but trending towards faulty, or in expected and/or acceptable operating condition.

[1758] In example embodiments, the operating characteristics detector 15020, the vision analytics module 15012, and/or the computer vision system 15000 may generate one or more warnings, signals, indicators, or other suitable outputs configured to alert the operator of one or more of the operating characteristics of the respective device 13006, of one or more components of the respective device 13006 that requires maintenance or replacement, any other suitable alert, or a combination thereof. For example, the computer vision system 15000 may be configured to generate a message, such as a text message, email message, popup message, or other suitable message, indicating that a component (e.g., the first pulley) of the respective device 13006 requires maintenance. The message may include text, characters, images, or other suitable information that

conveys the intended message. The computer vision system 15000 may be configured to communicate, via the network 15008, near field communication, or other suitable communication system or protocol, the message to the operator. For example, the computer vision system 15000 may communicate the message to a mobile device, as described, or other suitable device and/or location.

[1759] In example embodiments, the computer vision system 15000 may be configured to display on an output display a current status of one or more respective devices 13006. For example, a factory, plant, or other suitable location of the respective devices 13006 may include an output display (e.g., a screen or monitor) located such that operators within proximity of the respective devices 13006 can see the output display. The computer vision system 15000 may be configured to display a status (e.g., a red, yellow, green status, an up or down status, or other suitable status or indicator, or a combination thereof) of one or more of the respective devices 13006. For example, the computer vision system 15000 may display a green status next to the respective device 13006 that is operating within tolerable operating conditions (e.g., based on the visual analysis of the image data sets described above). In another example, the computer vision system 15000 may display a yellow status next to the respective device 13006 that is operating within tolerable operating conditions and the visual analysis indicates that the respective device 13006 may start to operated outside of the tolerable operating conditions if the operating characteristics (e.g., identified, as described) continue along a current operating trend (e.g., based on the frequency of vibration of the belt, the computer vision system 15000 determines that continued vibration at that frequency and/or increased frequency may cause the respective device 13006 to operate outside of the tolerable operating conditions). In another example, the computer vision system 15000 may display a red status next to the respective device 13006 that is currently operating outside of tolerable operating conditions. In example embodiments, the computer vision system 15000 may display the operating status of the respective devices 13006 on other suitable displays, such as a display of a mobile device, as described. For example, the mobile device may include an application that displays the operating status of the respective devices 13006.

[1760] In example embodiments, the output of the vision analytics module 15012 may be used to update and/or improve the training data sets, described above. For example, output from the vision analytics module 15012 may be used to update the training data sets to include additional operating characteristics, improve the precision of the values used to predict various operating characteristics, used for other suitable updates or improvements to the training data sets, or a combination thereof. The training data sets may be used as a continuous feedback to the computer vision system 15000 to improve predictive and determinative capabilities of the computer vision system 15000.

[1761] In example embodiments, the output of the vision analytics module 15012 may be used to populate and/or update a knowledgebase that may be used by an operator or by the computer vision system 15000 to identify faults, schedule repairs or maintenance, adjust settings on the respective devices 13006, take other corrective action, or other suitable action. For example, the output of the vision analytics module 15012 may be correlated with a corresponding repair of a component (e.g.,

the output of the vision analytics module 15012 may indicate that vibration of the belt is beyond the expected or acceptable tolerance and an operator may have replaced a pulley in response to the output). The knowledgebase may be updated to indicate that the output of the vision analytics module 15012 (e.g., including the values of the operating characteristics determined above) resulted in a replaced pulley. In this manner, the knowledgebase may continue to grow and provide accurate and precise information for an operator or the computer vision system 15000 as it relates to operating characteristics and corresponding corrective actions, thereby improving the efficiency of the computer vision system 15000 and assisting the operator in identifying issues and corresponding corrective actions.

[1762] In example embodiments, the computer vision system 15000 may be configured to visually inspect components, parts, systems, devices, or a combination thereof, other than those described above. For example, the computer vision system 15000 may be configured to visually inspect, as described, parts manufactured in a parts manufacturing facility. For example, the data capture devices 15002 may be disposed or positioned such that field of data capture for each respective data capture device 15002 is directed toward at least a portion of a part being manufactured (e.g., on a parts manufacturing line). The data capture devices 15002 may capture data associated with the parts as the parts move along the parts manufacturing line. The computer vision system 15000 may analyze the data captured by the data capture devices 15002 (e.g., as image data sets generated by the image data set generator 15006) and identify anomalies, variations, or other conditions that deviate from tolerable standards for the part. In example embodiments, the part may include a part for a vehicle, a part for a bike, a bike chain, a gasket, a fastener (e.g., a screw, a bolt, a nut, a nail, and the like), a printed circuit board, a capacitor, an inductor, a resistor, or other suitable part. For example, the computer vision system 15000 may analyze image data sets associated with bike chains being manufactured. The computer vision system 15000 may identify a bend in a portion of a bike chain that is outside of the tolerable standards for the portion of the bike chain based on the analysis described above. The computer vision system 15000 may generate a message, as described, indicating that the bike chain should be taken out of circulation, repaired, destroyed, or other suitable action.

[1763] As is generally illustrated in FIGS. 175-176, a deep learning system 15030 may be configured to train the computer vision system 15000, using the training data sets, to identify operating characteristics of the respective devices 13006 or other suitable devices, identify corrective actions in response to the identified operating characteristics, and initiate corrective action based on the identified corrective actions. The deep learning system 15030 may train the computer vision system 15000 using learning based on data representations. In example embodiments, the deep learning system 15030 may train the computer vision system 15000 using supervised training (e.g., using classification), semi-supervised training, or unsupervised training (e.g., using pattern analysis). In example embodiments, the deep learning system 15030 may include a deep neural network, a deep belief network, a recurrent neural network, other suitable networks or learning systems, or a combination thereof.

[1764] In example embodiments, the deep learning system 15030 may include propositional formulas or latent variables organized into a plurality of layers. Each of the plurality of layers may be configured to represent an abstract portion of an image. For example, a first layer may represent an abstract of pixels and encode edges of an input image, for example, an image representing a point of interest of the representative device 13006. A second layer may represent arrangements of the edges. A third layer may encode a first portion of a component within the point of interest of the representative device 13006 (e.g., a portion of the belt, as described). A fourth later may represent another encoded portion of the component, and so on, such that, the plurality of layers, when overlaid, represents the point of interest of the representative device 13006. The deep learning system 15030 may be configured to translate the layers into training data sets, used to train the computer vision system 15000. For example, the deep learning system 15030 may translate a plurality of layers of one or more images that represent a belt of the representative device 13006 vibrating at a first frequency. The deep learning system 15030 may use input data from various sources to determine whether the first frequency represents a frequency at which the belt is vibrating within the expected or acceptable tolerances, as described. For example, the deep learning system 15030 may receive data indicating repair data, maintenance data, uptime data, downtime data, profitability data, efficiencies data, operational optimization data, other suitable data, or a combination thereof, associated with the respective device 13006, a process, a production line, a facility, or other suitable systems.

[1765] In example embodiments, the deep learning system 15030 may identify data values corresponding to the first frequency of the belt. For example, the deep learning system 15030 may identify an uptime value, a downtime value, a profitability value, other suitable values, or a combination thereof that correspond to periods when the respective device 13006 operated with the belt vibrating at the first frequency. For example, the deep learning system 15030 may determine that the first frequency is within the expected or acceptable tolerances when the data indicates that the respective device 13006 had an uptime that was above a threshold, a downtime that was below a threshold, a profitability that was above a threshold, or a combination thereof. Conversely, the deep learning system 15030 may determine that the first frequency is beyond the expected or acceptable tolerances when, for example, the downtime associated with the respective device 13006 was above a threshold. It should be understood that the deep learning system 15030 may identify any suitable operating characteristic besides those disclosed herein and that the deep learning system 15030 may determine positive or negative outcomes of the operating characteristics based on any suitable data analysis other than those described herein.

[1766] In example embodiments, the deep learning system 15030 may generate the training data sets using the identified operating characteristics and associated analysis thereof. In example embodiments, the deep learning system 15030 may train the computer vision system 15000 using the training data sets. In example embodiments, the deep learning system 15030 may receive feedback information from the computer vision system 15000, an operator, a programmer, other suitable sources, or a combination thereof. The deep learning system 15030 may update the training data sets based on the feedback. For example, the computer vision system 15000, having been

trained using the training data sets, may identify a component as faulty. The operator may visually inspect the component and determine that the component is not faulty. The operator and/or the computer vision system 15000 may communicate to the deep learning system 15030 that the component was not faulty based on the identified operating characteristics (e.g., identified by the computer vision system 15000). The deep learning system 15030 may update the training data sets using the feedback from the operator and/or the computer vision system 15000.

[1767] In example embodiments, a computer vision system for detecting operating characteristics of a manufacturing device, includes at least one data capture device configured to capture raw data of a point of interest of the manufacturing device, a memory, and a processor. The memory includes instructions executable by the processor to: generate one or more image data sets using the raw data captured; visually identify one or more values corresponding to a portion of the manufacturing device within the point of interest represented by the one or more image data sets; record the one or more values; visually compare the recorded one or more values to corresponding predicted values; generate a variance data set based on the comparison of the recorded on or more values and the corresponding predicted values; identify an operating characteristic of the manufacturing device based on the variance data; compare the operating characteristic to a threshold; determine whether the operating characteristic is within a tolerance based on whether the operating characteristic is greater than the threshold; and generate an indication indicating the operating characteristic.

[1768] In example embodiments, the computer vision system is trained by a deep learning system. In example embodiments, the deep learning system is configured to train the computer vision system using at least one training data set. In example embodiments, the at least one training data set includes image data. In example embodiments, the at least one training data set includes non-image data.

[1769] In example embodiments, a computer vision system for detecting operating characteristics of a device, includes at least one data capture device configured to capture raw data of a point of interest of the device, a memory and a processor. The memory includes instructions executable by the processor to: generate one or more image data sets using the raw data captured; visually identify one or more values corresponding to a portion of the device within the point of interest represented by the one or more image data sets; record the one or more values; visually compare the recorded one or more values to corresponding predicted values; generate a variance data set based on the comparison of the recorded one or more values and the corresponding predicted values; identify an operating characteristic of the device based on the variance data; compare the operating characteristic to a threshold; determine whether the operating characteristic is within a tolerance based on whether the operating characteristic is greater than the threshold; and generate an indication indicating the operating characteristic.

[1770] In example embodiments, the device includes an agitator. In example embodiments, the device includes an airframe control surface vibration device. In example embodiments, the device includes a catalytic reactor. In example embodiments, the device includes a compressor. In example embodiments, the device includes a conveyor. In example embodiments, the device

includes a lifter. In example embodiments, the device includes a pipeline. In example embodiments, the device includes an electric powertrain. In example embodiments, the device includes a robotic assembly device. In example embodiments, the device includes a device in a gas production environment. In example embodiments, the device includes a device in a pharmaceutical environment.

[1771] In example embodiments, flow of information among participants and elements of a predictive maintenance knowledge platform may be configured as depicted in Figure 177. A platform 28600 as exemplary configured in Figure 177 may include a plurality of subsystems that may include one or more of: data storage, machine intelligence, and industrial machine-related transactions. Such a subsystem may be a web-server based system, a distributed system, a handheld device, an industrial machine co-resident system, and the like. In an example, the industrial machine maintenance data analysis subsystem 28602 may include a data storage 28604, machine learning and/or an artificial intelligence facilities 28606, a transaction facility 28608 and the like. The Industrial machine maintenance data analysis subsystem 28602 may provide services 28610 including updates to industrial machine related data, such as service criteria, fault prevention, service pricing, parts pricing, tests and criteria for detecting potential machine faults, analysis of repairs and the like, functions and updates to fault prediction metadata, and the like. The industrial machine maintenance data analysis subsystem 28602 may provide information, such as those associated with the provided services 28610, in the form of streams, transactions, data base reading and writing, and the like for access to cloud-based data storage. The industrial machine maintenance data analysis subsystem 28602 may receive information regarding individual industrial machines from the machines via the data collection network 28612. In example embodiments, a data collection network 28612 may be described herein and in the documents referenced and incorporated herein. The industrial machine maintenance data analysis subsystem 28602 may receive information from specific industrial machines such as machine parameters and the like that may be retrieved from one or more smart RFID elements 28614 of the industrial machine. In example embodiments, smart RFID elements may be configured with portions of industrial machine and may have functionality as described elsewhere herein.

[1772] In example embodiments, an industrial machine predictive maintenance subsystem 28616 may apply machinery fault detection, identification, classification, and related algorithms to the data provided from the industrial machine maintenance data analysis subsystem 28602 and to data further provided from an industrial machine health monitoring facilities 28618 and the like to generate data structures, streams, and other electronic data that may be communicated to facilitate predictive maintenance of industrial machines. In example embodiments, the industrial machine predictive maintenance subsystem 28616 may receive and analyze a stream or the like of industrial health monitoring data from the industrial machine health monitoring facility 28618. One or more results of such stream analysis may include determination of conditions that indicate a healthy machine, an unhealthy machine, a likelihood of at least a portion of a machine that may need service to avoid a fault, a specific machine that requires service, and the like. Conditions that may indicate a healthy machine may be a result of tests and the like performed on or by industrial

machines and communicated to the machine health monitoring facility 28618. In an example, the machine health monitoring facility 28618 may receive operation-related information, such as sensor data from industrial machine motors (e.g., torque, revolutions per minute, run time, start/stop data, directional data and the like) in a live or delayed stream from one or more industrial machines. This operation-related data may be processed by the health monitoring facility 28618 to detect when, for example, a number of revolutions over a set period of time, such as a day, week, month and the like exceeds a maintenance threshold value. A portion of the stream data and/or the result of processing by the health monitoring facility 28618 may be provided, such as a stream and the like to the industrial machine predictive maintenance subsystem 28616 for uses as described, including identifying potential faults and the like that are to be addressed with predictive maintenance and the like. The industrial machine predictive maintenance subsystem 28616 may generate one or more predictive maintenance sets of data 28620 that may identify one or more industrial machines and may indicate portion(s) of the machine that are determined to benefit from service, maintenance, repair, replacement and the like. The sets of data 28620 may include specific parts, service procedures, materials, service timeframes, required to perform a predictive maintenance activity on one or more specific industrial machines. In example embodiments, machine fault analysis that may be performed by the industrial machine predictive maintenance subsystem 28616 may facilitate generating work orders from a CMMS subsystem 28622.

[1773] In example embodiments, the CMMS subsystem 28622 may receive industrial machine details, service (e.g., repair, maintenance, upgrade, and the like) details for the industrial machine, procedures to be followed, parts needed, and the like from sources such as the industrial machine predictive maintenance subsystem 28616, a CMMS interface 28624, data structures configured and maintained that may include parts lists and the like for the industrial machine and any other information to facilitate performing service on the industrial machine. The CMMS subsystem 28622 may initiate actions with parts suppliers, service providers, third-party partners, vendors, an owner/operator of the industrial machine to be serviced and the like. In an example, the CMMS subsystem 28622 may generate orders for services from one or more service providers that are known to the CMMS subsystem 28622 as qualified to provide the services required.

[1774] In example embodiments, the CMMS subsystem 28622 may interface with one or more predictive maintenance knowledge bases and/or knowledge graphs that may be stored in a data store accessible by the CMMS subsystem. In example embodiments, such a CMMS knowledge base or the like may further include a knowledge graph that may contain information beneficial to the service determination and order generation services provided by the CMMS subsystem 28622. A CMMS knowledge graph may contain or provide computer access to information about industrial machines, service activity of industrial machines, costs (e.g., historical, trending, and predictive) for parts, materials, tools, and services of industrial machines, algorithms and functionality for delivering the CMMS services 28626 and the like. The CMMS subsystem 28622 may facilitate coordination with service providers, parts providers, material and tool providers and the like based on an industrial machine owner's decision regarding servicing the industrial machine so that the service can be performed in a timeframe that the owner chooses.

[1775] The CMMS subsystem 28622 may access information in the smart RFID element(s) 28614 via the CMMS interface 28624 that may facilitate access to individual industrial machines and the like. The CMMS subsystem 28622 may use information received via the CMMS interface 28624 to facilitate performing coordination of resources to perform maintenance effectively and efficiently for the specific machine. In an example, a specific industrial machine may have an operating cycle that results in greater utilization of one of its moving parts (e.g., an industrial motor) than typical. This information may be processed by the predictive maintenance subsystem 28616 and result in an indication of a service that may need to be performed on the machine. The predictive maintenance subsystem 28616 may provide information to the CMMS subsystem 28622 that it would process to generate orders for parts, services, and the like. This knowledge may be used by the CMMS subsystem 28622 to interact with service, parts, and material suppliers to provide a firm quote for performing a utilization-based maintenance service at a different time (e.g., weeks or months sooner) than other comparable industrial machines with lower utilization rates.

[1776] In example embodiments, the CMMS subsystem 28622 may execute algorithms that gather information about a plurality of industrial machines, including a plurality of industrial machines of different types of machine (e.g., stationary machines, mobile machines, machines on vehicles, machines deployed at job sites, and the like) along with service provider information, parts and parts provider information, part location and inventory information, machine production providers, third-party parts handlers, logistics providers, transportation providers, service standards, service requirements, service activities including results of service and the like, and other information to facilitate providing services 28626 including coordinating orders for services, parts and the like.

[1777] In example embodiments, in response to industrial machine fault identification information provided from the preventive maintenance subsystem 28616, the predictive maintenance knowledge system 30002 may identify candidate service providers. Service providers that are known to the CMMS subsystem 28622 as having successfully demonstrated experience with the procedure needed for the requested service may be contacted to provide a service estimate and/or a price estimate for service, parts, and the like. Similarly, parts and/or material that may be associated with the procedure of the requested service may be identified. Factors such as part cost, transportation costs, availability, location of the parts versus the machines, prior relationships between one or more parts providers and a party associated with the service request, such as the industrial machine owner and the like, and other factors may be evaluated to determine which parts provider to contact in preparation for ordering the parts. With these factors considered, a part inquiry may be placed with one or more parts providers in anticipation of the service being conducted by the qualified service indication from the preventive maintenance subsystem 28616 with one or more service recommendations. In example embodiments, the CMMS subsystem 28622 may have enough information to automatically select a specific service recommendation and may, with or without explicit approval, generate a service order 28626 that may include a parts/material/tools order if needed for the requested service.

[1778] In example embodiments, information that the CMMS subsystem 28622 may rely on may be sourced from an Enterprise Resource Planning (ERP) interface associated with the industrial machine as well as third-party sources of information such as independent parts suppliers, service providers, and the like that may offer parts and/or services for industrial machines. In example embodiments, the CMMS subsystem 28622 may coordinate with an industrial machine owner's ERP system, such as via the ERP interface 28628 to effect placement of orders with the service provider, parts provider, and the like. The CMMS subsystem 28622 may use service material provider information to determine price and availability of service material. This information may be combined with service material inventory information to facilitate generating suitable orders for service material as part of the industrial machine service offering 28626.

[1779] In example embodiments, the CMMS subsystem 28622 may receive a timeframe in which the repair must be completed in order to avoid failure and the recommended repair with instructions from the manufacturers manual on how to conduct the repair. This repair information may be then processed by the CMMS subsystem 28622 (e.g., a cloud based system) where a work order is created and tracked. The work order may be digitally pushed to the ERP system to check the plant's production schedule to find when the specific machine requiring maintenance is available for repair based on the time frame provided by the analysis and the amount of time the machine will be off-line based on, for example information in a manufacturer's manual referenced in a service procedure that states how much time it should take to make the repair. Once the ERP system finds the available date it may coordinate with the CMMS subsystem 28622 to ask for bids from vendors for the parts and the service work or to place orders for the parts and with a service contractor, such as a preferred contractor. In example embodiments, the CMMS subsystem 28622 or the ERP system may configure a request for bids by simply using the manufacturers manual for the procedure to provide the bidders with the required parts information (e.g., part numbers, vintage, revision, specifications, after-market alternatives, last price paid, if a used part is OK, and the like) and the repair actions necessary for the service action (e.g., the procedure steps, diagnostics, equipment/tools required, materials required, personnel required, and the like). A bid may be based on the repair actions listed in the procedure and may become the scope of work for the job to be bid on. In example embodiments, if there are other problems found and addressed outside of this scope a secondary process may be followed to approve additional compensation to the vendor.

[1780] In example embodiments, a service delivery and tracking subsystem 28630 may be used by service providers, such as service technicians, industrial machine owners/operators, third parties (e.g., auditors, regulators, union personnel, safety associations, parts manufacturers and the like) to gather and report information associated with an ordered service request as may be determined from service order data 28626. The service delivery and tracking subsystem 28630 may include functionality that matches up machine procedures with service requirements, ensures that images associated with the ordered service (e.g., a part being services, an installation of the machine, a video of the machine operating before and/or after service, parts that have been removed from the industrial machine, service personnel, and the like) are captured with sufficient quality to meet image quality standards for automatic detection of one or more parts of the industrial machine.

[1781] In example embodiments, the service delivery and tracking subsystem 28630 may report data, repairs, images and the like, collectively service data 28632 to an industrial machine maintenance data analysis subsystem 28602 for refinement of service procedures, parts ordering, and the like.

[1782] In example embodiments, compensation for work and analysis performed by the various subsystems may be derived from various sources. The CMMS subsystem 28622 operator / owner / affiliate may be compensated on a transaction basis, such as by receiving a fee for each part or service ordered. Such a fee may include a fixed portion (e.g., amount per part order) and may include a variable portion (e.g., a percent of an order total). This fee may be explicitly included in charges billed to a party responsible for payment of the parts and services to perform the maintenance action. This fee may be built into the cost of each part/service and recovered as a deduction from the payment that is passed from the responsible party to the parts and/or service provider.

[1783] In example embodiments, an industrial machine predictive maintenance system may include an industrial machine data analysis facility that generates streams of industrial machine health monitoring data by applying machine learning to data representative of conditions of portions of industrial machines received via a data collection network. The system may further include an industrial machine predictive maintenance facility that produces industrial machine service recommendations responsive to the health monitoring data by applying machine fault detection and classification algorithms thereto. The system may further include a computerized maintenance management system (CMMS) that produces at least one of orders and requests for service and parts responsive to receiving the industrial machine service recommendations. And, the system may include a service and delivery coordination facility that receives and processes information regarding services performed on industrial machines responsive to the at least one of orders and requests for service and parts, thereby validating the services performed while producing a ledger of service activity and results for individual industrial machines.

[1784] In example embodiments, methods and systems for finding a set of workers having relevant know-how and expertise about maintenance, service and repair of a specific machine may employ machine learning algorithms with worker selection algorithms to ensure timely, quality workers are selected and deployed for industrial machine servicing, such as for predictive maintenance and the like described herein. Referring to Figure 178, machine learning-based methods 32400 for finding a set of workers as described above is depicted. In example embodiments, the facility for finding workers 32402 may be configured as a system that may include a set of algorithms and data structures that may execute on a processor. The worker finding facility 32402 may process data about workers, machines, procedures, and the like with algorithms that facilitate matching qualified workers with service activities, such as predictive maintenance activities and the like. In an example of finding workers, a service activity may include following a service or maintenance procedure 32406, such as to repair and/or maintain a portion of an industrial machine. The procedure 32406 may further indicate one or more industrial machines, such as by model number, family, and the like. The worker finding facility 32402 may further access, such as by retrieving

information about workers from a worker database 32422, information that facilitates characterizing one or more workers, including procedures for which the worker has experience, training, certification and the like. One or more workers who have experience and the like with the procedure may be selected for further refinement, which may include matching a worker location to a machine location, a worker availability and/or schedule to a machine service schedule, worker rates/fees to machine owner service budgets and the like. One or more workers on a resulting list of refined workers may be contacted about a service to be performed on the machine. Based on, for example, replies to such worker contact, a primary worker may be selected by the worker finding facility 32402 and allocated to perform the service via the procedure 32406.

[1785] In example embodiments, the worker finding facility 32402 may access a list of procedures 3246 for which service may be required. The worker finding facility 32402 may build a data set of workers that qualify for performing the procedure, such as by searching through worker information 32416 for workers who meet procedure criteria, such as a number of times the worker has performed the procedure, a number of times a worker has performed a similar procedure, and the like. Workers with more experience may be marked as preferred workers in such a database for the specific procedure so that when the procedure is required to be performed, those preferred workers may be readily identified. In example embodiments, workers may directly maintain the worker database 32422 by updating information regarding procedures and the like that they perform.

[1786] In example embodiments, the worker finding facility 32402 may receive information about procedures 32406, machines 32408, machine location 32410, machine owner and/or affiliation 32412, required service schedule 32414 and the like for one or more service activities, such as a predictive maintenance activity and the like to be performed and form a profile of a preferred worker for a given combination of procedure, machine, location, owner, schedule and the like. The worker finding facility 32402 may build a profile for various combinations of such information so that workers that best meet the profile may be readily found. In example embodiments, such preferred worker profiles may be published so that third parties, such as service organizations and the like may provide estimates and the like for providing a service based on the profile. These estimates may be captured and used by the methods and systems of predictive maintenance of industrial machines and the like to build a marketplace of service providers for common or often required services, such as preventive maintenance services and the like.

[1787] In example embodiments, information captured in the worker database 32422 and the like may be processed with machine learning algorithms 32424 to facilitate improving matching of workers with requirements for providing qualified workers for procedures and the like. In example embodiments, the preferred worker profiles and information received in response to their publication may be processed with the machine learning algorithms 32424 to refine the algorithms that are used to build preferred worker profiles.

[1788] In example embodiments, additional information that may influence worker selection by the worker finding facility 32402 may include affiliation of the worker with service organizations, manufacturers of industrial machines, industry organizations, and the like. Referrals and or

feedback on specific workers may be factored into determination of individual workers, worker groups and the like as to their preferred worker status and the like. Worker rates and/or fees (e.g., based on estimates, actual charges, payment terms and the like) may further be factored into finding a worker, such that workers that when two or more workers overall have comparable qualifications, a worker with lower costs or easier payment terms may be ranked higher for a given procedure than one with higher cost and the like.

[1789] In example embodiments, techniques for finding workers may be performed in real-time or near real time as demands for industrial machines require. In this way, as new workers become available, finding a worker may incorporate updates to worker profiles and the like that may be accessible over websites, and the like via the Internet.

[1790] In example embodiments, a system may include an industrial machine predictive maintenance facility that produces industrial machine service recommendations by applying machine fault detection and classification algorithms to industrial machine health monitoring data. Such a system may also include a worker finding facility that identifies at least one candidate worker for performing a service indicated by the industrial machine service recommendations by correlating information in the recommendation regarding at least one service to be performed with at least one of experience and know-how for industrial service workers in an industrial service worker database. In example embodiments, the system may include machine learning algorithms executing on a processor that improve the correlating based on service-related information for a plurality of services performed on similar industrial machines and worker-related information for a plurality of services performed by the at least one candidate worker.

[1791] In example embodiments, an industrial machine maintenance part/service ordering facility 32502 for industrial machine service and maintenance 32500, including predictive maintenance and the like may be embodied as depicted at least in Figure 179 filed herewith. The industrial machine maintenance part/service ordering facility 32502 may facilitate finding, ordering, and fulfilling orders for relevant parts and components, so that maintenance, service and repair operations for industrial machines can occur seamlessly, with minimal disruption. In example embodiments, the industrial machine maintenance part/service ordering facility 32502 may receive industrial machine details 32508, service (e.g., repair, maintenance, upgrade, and the like) details 32510 for an industrial machine, procedures to be followed 32506, parts needed 32514, service providers 32520, parts providers 32522 and the like. The industrial machine maintenance part/service ordering facility 32502 may initiate actions with parts suppliers, service providers, third-party partners, vendors, an owner/operator of the industrial machine to be serviced and the like. In an example, the industrial machine maintenance part/service ordering facility 32502 may generate orders for services 32518 from one or more service providers 32520 that are known to the industrial machine maintenance part/service ordering facility 32502 as qualified to provide the services required. The industrial machine maintenance part/service ordering facility 32502 may also generate orders for parts 32516 from one or more parts providers 32522 that are known as qualified to provide the parts required, on time, within budget, and the like. The parts orders 32516 and the service orders 32518 may also be communicated to an owner 32512 or other entity

responsible for ensuring access to the industrial machine. The parts and service providers selected may further coordinate with the owner 32512 to ensure the service can be properly delivered. The industrial machine maintenance part/service ordering facility 32502 may have access to the machine owner 32512 preferences and/or requirements regarding scheduling, budgets, service and parts provider preferences and/or affiliations, and the like to facilitate coordination with service providers, parts providers, material and tool providers and the like based thereon.

[1792] Factors such as part cost, transportation costs, availability, location of the parts versus the machines, prior relationships between one or more parts providers and a party associated with the service request, such as the industrial machine owner and the like, and other factors may be evaluated to determine which parts provider 32522 to contact in preparation for ordering the parts 32516. With these factors considered, a part inquiry may be placed with one or more parts providers 32522 in anticipation of the service being conducted by the qualified service provider. In example embodiments, the industrial machine maintenance parts/service ordering facility 32502 may have enough information to automatically select a specific service provider 32520 and may, with or without explicit approval, generate the service order 32518.

[1793] In example embodiments, information that the industrial machine maintenance part/service ordering facility 32502 may rely on information regarding vendors, and the like from an Enterprise Resource Planning (ERP) system owned and or operated by the owner of the industrial machine. In example embodiments, the industrial machine maintenance part/service ordering facility 32502 may coordinate with an industrial machine owner's ERP system to effect placement of orders with the service provider, parts provider, and the like.

[1794] In example embodiments, a system may include an industrial machine maintenance part and service ordering facility that prepares and controls orders for parts and services responsive to service recommendations received from an industrial machine predictive maintenance facility that produces industrial machine service recommendations by applying machine fault detection and classification algorithms to industrial machine health monitoring data. In example embodiments, the system may further analyze a procedure associated with the service recommendations for generating at least one of the orders for parts and services.

[1795] In example embodiments, an industrial machine predictive maintenance system may include deployment of smart RFID devices on portions of industrial machines. The smart RFID devices may be configured to include information about the machine, such as configuration information, assembly information, physical element details (e.g., part numbers, revisions, production details, test details, and the like), procedure information (e.g., assembly, disassembly, test, configuration, service, parts replacement, and the like), and other operational information and the like. Smart RFID devices may be disposed with each major element in a machine, such as each element that might include information relevant for efficient service and maintenance of the machine. In example embodiments, disposing smart RFID devices may be configured into the production of industrial machine and the like parts and sub systems so that production information and the like of the part(s) can be captured for the specific part, and the like. A smart RFID element may not only provide storage for a range of information, including large service manuals and the

like, a smart RFID element may include functionality, such as searching, indexing, linking, and the like that may facilitate users quickly finding procedures, such as lubricating procedures, bearing replacement procedures, bearing fault frequencies, and the like that may be crucial for machine trouble shooting and the like. In example embodiments, at least one method for accessing the information may be compatible with existing techniques used by expert service personnel, which may be taught to new service providers while these experts remain on the job. In example embodiments, providing easy access, including indexing, linking and the like may be built into the documents, procedures, data sheets, manuals and the like during their creation so that common access approaches can be used for any embodiment of the information (e.g., in the smart RFID, in a cloud representation of the RFID, in 3rd party service manuals, in industrial machine producer systems and the like).

[1796] Referring to Figure 180, an industrial machine 32600 may be configured from a plurality of elements, parts, sub-assemblies and the like. One such sub-assembly might include an industrial machine motor 32602. An RFID device may be disposed with the machine that may include details, such as those described herein for smart RFID devices, for the specific motor. The motor 32602 RFID device may communicate, such as through wireless communication with other devices brought into proximity, such as a smart phone, tablet or the like 32614 so that a user of the table and the like 32614 may access the information stored on the motor 32602 RFID device for conducting service, maintenance, testing, and the like. In example embodiments, the motor 32602 service procedure may be retrieved from the motor 32602 RFID and displayed via an application executing on the table 32614 to be followed by the service technician. Another such sub-assembly might include an industrial machine drive shaft 32604. An RFID device may be disposed with the machine that may include details, such as those described herein for smart RFID devices, for the specific drive shaft 32604. The drive shaft 32604 RFID device may communicate, such as through wireless communication with other devices brought into proximity, such as a smart phone, tablet or the like 32614 so that a user of the table and the like 32614 may access the information stored on the drive shaft 32604 RFID device for conducting service, maintenance, testing, and the like. In example embodiments, the drive shaft 32604 service procedure may be retrieved from the drive shaft 32604 RFID and displayed via an application executing on the table 32614 to be followed by the service technician. Yet another such sub-assembly might include an industrial machine gear box 32606. An RFID device may be disposed with the machine that may include details, such as those described herein for smart RFID devices, for the specific gear box 32606. The RFID device in the gear box 32606 device may communicate, such as through wireless communication with other devices brought into proximity, such as a smart phone, tablet or the like 32614 so that a user of the table and the like 32614 may access the information stored on the gear box 32606 RFID device for conducting service, maintenance, testing, and the like. In example embodiments, the gear box 32606 service procedure may be retrieved from the gear box 32606 RFID and displayed via an application executing on the table 32614 to be followed by the service technician. Yet another such sub-assembly might include an industrial machine articulated arm 32608. An RFID device may be disposed with the machine that may include details, such as those described herein

for smart RFID devices, for the specific articulated arm 32608. The articulated arm 32608 RFID device may communicate, such as through wireless communication with other devices brought into proximity, such as a smart phone, tablet or the like 32614 so that a user of the table and the like 32614 may access the information stored on the articulated arm 32608 RFID device for conducting service, maintenance, testing, and the like. In example embodiments, the articulated arm 32608 service procedure may be retrieved from the articulated arm 32608 RFID and displayed via an application executing on the table 32614 to be followed by the service technician.

[1797] Referring further to Figure 180, yet another such sub-assembly might include an industrial machine bucket 32610. An RFID device may be disposed with the machine that may include details, such as those described herein for smart RFID devices, for the specific bucket 32610. The bucket 32610 RFID device may communicate, such as through wireless communication with other devices brought into proximity, such as a smart phone, tablet or the like 32614 so that a user of the table and the like 32614 may access the information stored on the bucket 32610 RFID device for conducting service, maintenance, testing, and the like. In example embodiments, another such sub-assembly might include an industrial machine drive train 32612. An RFID device may be disposed with the machine that may include details, such as those described herein for smart RFID devices, for the specific drive train 32612. The drive train 32612 RFID device may communicate, such as through wireless communication with other devices brought into proximity, such as a smart phone, tablet or the like 32614 so that a user of the table and the like 32614 may access the information stored on the drive train 32612 RFID device for conducting service, maintenance, testing, and the like. In example embodiments, the drive train 32612 service procedure may be retrieved from the drive train 32612 RFID and displayed via an application executing on the table 32614 to be followed by the service technician. In example embodiments, any of the RFID devices, such as the motor 32602 RFID, the drive shaft 32604 RFID, the gear box 32606 RFID, the articulated arm 32608 RFID, the bucket 32610 RFID, the drive train 32612 RFID and the like may communicate via a wireless communication network with an access point, such as industrial machine access point 32616 that may be disposed on the industrial machine 32600 or proximal thereto. Communication from the RFID devices through the industrial machine access point 32616 to gain access to a network 32618, such as a network for connecting other industrial machines in a facility or external networks such as the Internet. Information stored in the industrial machine RFID devices may be transmitted over the network 32618 for use in the predictive maintenance methods and systems described herein.

[1798] In example embodiments, a system may include a smart RFID element configured to capture and store in a non-volatile computer-accessible memory operational, physical and diagnostic result information for a portion of an industrial machine by communicatively coupling with at least one sensor configured to monitor a condition of the portion of the industrial machine. The smart RFID element may further be configured to receive, organize, and store in the non-volatile memory information that enables execution of at least one service procedure for the industrial machine.

[1799] In example embodiments, information about an industrial machine, such as about a portion of the industrial machine may be stored in an RFID element disposed with the industrial machine or portion thereof. The information stored may be configured to facilitate rapid random access to any portion of the information quickly and efficiently, such as through use of a smart phone or other computing device configured with at least a web browser and the like. The information may be configured as one or more data structures, such as a hierarchical data structure and the like that may also facilitate exploration of the information through browsing the hierarchy and the like. Referring to Figure 181, an exemplary high level structure 32700 of a portion of such an RFID is presented and includes rows and columns. The exemplary high level structure 32700 may include a category of information 32702 that may identify a general area of information, such as production and the like. Each such category may be described in a description column 32704 that may have further identifying information. A notes column 32706 may be configured with free-form notes that may be updated as needed. In example embodiments, the category 32702 may include a range of information categories associated with the industrial machine, such as Production, Parts, Quality, Installation, Validation, Procedures, Operational, Assembly and the like. In an example of the category 32702, validation 32708 may include a list of validation tests that are required and that are performed, along with results. Validation tests may be performed to validate installation at a customer site and the like. Validation 32708 may also include links to one or more procedures accessible in the RFID through the procedures 32710 category that are required for validation.

[1800] In example embodiments, industrial machine-related information that may be stored on and/or accessible via a smart RFID element may include, without limitation operational data collected by sensors deployed with the industrial machine and collected via the sensor data collection methods and systems described and the references incorporated herein. Other information that may be stored on or accessible from a smart RFID element may include, without limitation detected exceptions in operational and/or test data, such as excess temperatures, unexpected shutdowns, system restarts, and the like. A smart RFID element may communicate with an external computing device, such as a smart phone, tablet, communication infrastructure node, computer, mesh network device, and the like via a range of communication protocols including Wi-Fi, NFC, BLUETOOTH™, and others. In example embodiments, a smart RFID element may communicate wirelessly with a portable computing device when the computing device is in wireless communication proximity, such as when a portable computing device is brought within NFC range of the smart RFID element. A smart RFID element may communicate over a network, such as the Internet as an IoT device. The smart RFID element may send data to a server, such as a web server or the like that may aggregate information from the element and cloud-accessible sources for one or more service activities associated with the industrial machine. In example embodiments, a smart RFID element may communicate with external computing device(s) at convenient times, such as at the end/start of an activity, shift, day, when preventive maintenance is soon to be performed, and the like.

[1801] A smart RFID element may be used during production and/or assembly of an industrial machine or portion thereof to capture physical details of the machine, such as for bearing

frequency, gear teeth count and type, build/assembly version information, build/test parameters, self-test information, calibration information, test time, inventory dwell time, and the like.

[1802] A smart RFID element may be used during installation and/or deployment of an industrial machine or portion thereof to capture orientation of the machine, testing activity, start-up activity, validation activity/runs, production start time, installation/deployment/configuration personnel, images of the industrial machine, and the like, at least a portion of which may be determined by one or more installation and/or deployment procedures that may be stored on and/or accessible through the smart RFID element.

[1803] In example embodiments, a system may include a smart RFID element configured to capture and store in a non-volatile computer-accessible memory operational, physical and diagnostic result information for a portion of an industrial machine by communicatively coupling with at least one sensor configured to monitor a condition of the portion of the industrial machine. The smart RFID element may further be configured to receive, organize, and store in the non-volatile memory information that enables execution of at least one service procedure for the industrial machine. The smart RFID may further be configured to facilitate hierarchical access to information about the industrial machine, including a plurality of portions directly accessible from a root entry for the industrial machine. In example embodiments, each of the plurality of directly accessible portions is structured to store entries for one portion selected from the list consisting of production information, parts information, quality information, installation information, validation information, procedure information, operational information, and assembly information.

[1804] In example embodiments, an alternate configuration of a smart RFID for industrial machine information storage and access, such as for service and the like may include a data structure as depicted in Figure 182. Data structure 32800 may be organized as columns and rows as shown, and the like. A first column may be a topic column 32802, such as production topics including, without limitation, date(s) of assembly, location, model number, serial number, time, work order number, customer, images of the industrial machine as built and the like. Each topic in the topic column 32802 may have one or more corresponding values in a value column 32804. In an example, a serial number topic 32808 in the topic column 32802 may have one or more corresponding serial numbers for the specific industrial machine listed in the value column 32804. Comments or other meta data for each topic in the topic column 32802 may be captured in corresponding entries in a notes column 32810.

[1805] In example embodiments, a system may include a smart RFID element configured to capture and store in a non-volatile computer-accessible memory operational, physical and diagnostic result information for a portion of an industrial machine by communicatively coupling with at least one sensor configured to monitor a condition of the portion of the industrial machine. The smart RFID element may further be configured to receive, organize, and store in the non-volatile memory information that enables execution of at least one service procedure for the industrial machine. In example embodiments, the production portion may include entries for assembly date, assembly location, machine model number, machine serial number, machine

assembly time, machine assembly work order number, customer, and images of portions of the industrial machine.

[1806] In example embodiments, an alternate configuration of a smart RFID for industrial machine information storage and access, such as for service and the like may include a procedure data structure as depicted in Figure 183. A machine-level procedure data structure 32900 may be organized as columns and rows as shown, and the like. A first column may be a procedure column 32902 that may list machine-level procedures, such as calibration, shutdown, regulatory compliance, assembly, safety-checking, image capture and the like. Each procedure in the machine-level procedure column 32902 may have one or more corresponding values in an attribute column 32904, such as a procedure identification number, a version, and the like. In an example, a safety check procedure 32908 entry in the procedure column 32902 may have one or more corresponding procedure number(s) and corresponding version number(s) in the column 32904. Comments or other meta data for each procedure in the procedure column 32902 may be captured in corresponding entries in a notes column 32910.

[1807] In example embodiments, a system may include a smart RFID element configured to capture and store in a non-volatile computer-accessible memory operational, physical and diagnostic result information for a portion of an industrial machine by communicatively coupling with at least one sensor configured to monitor a condition of the portion of the industrial machine. The smart RFID element may further be configured to receive, organize, and store in the non-volatile memory information that enables execution of at least one service procedure for the industrial machine. In example embodiments, the procedure portion may include entries for procedures selected from the list consisting of calibration, shutdown, regulatory, assembly, safety check, image capture, preventive maintenance, part repair, part replacement, and disassembly.

[1808] In example embodiments, referring to FIG. 184, methods and systems for collecting information 33000 about an industrial machine 33020, such as information about the machine operation, conditions, and the like may be beneficial to industrial machine predictive maintenance methods and systems, such as those described herein and elsewhere. In example embodiments, collecting the information from sensors on an industrial machine may include routing the collected information through one or more access points 33008 to a networked server 33018 where the information may be processed and stored. In example embodiments, collecting information from sensors on an industrial machine may include communicating between sensors and a smart RFID device 33002 disposed on or with the machine. Data from sensors, such as temperature sensors 33010, vibration sensors 33012, rotation sensors 33014, operational cycle sensors (e.g., cycle counters and the like) 33016 may be provided to a smart RFID device 33002 where the information may be processed and stored for further access by an external device, such as the server 33018, a handled device (not shown) brought into communication proximity of the industrial machine 33020, and the like. Industrial machine-specific data may be collected from the sensors and routed to one or more web servers 33018 that may employ a processor 33006 to generate a digital twin 33004 of the smart RFID 33002 on a computer accessible memory other than the smart RFID 33002. In example embodiments, the digital twin 33004 may be generated by copying content in

the smart RFID 33002. Likewise, machine-specific sensed data may be copied from the RFID twin 33004 memory to the smart RFID device 33002. Therefore, the RFID twin 33004 may be a copy of the smart RFID 33002, may be created independently of the smart RFID 33002, while maintaining a compatible structure, format, and substantively identical content, or may be a source of machine-specific data (e.g., as provided from the sensors over the access point) that may be copied to the smart RFID 33002 to maintain a copy of the information on the machine. In example embodiments, server 33018 may maintain a digital twin of a plurality of smart RFID devices for a plurality of industrial machines, including multiple smart RFID devices for a single industrial machine and the like.

[1809] In example embodiments, a system may include a smart RFID element configured to capture and store in a non-volatile computer-accessible memory operational, physical and diagnostic result information for a portion of an industrial machine by communicatively coupling with at least one sensor configured to monitor a condition of the portion of the industrial machine. The smart RFID element may further be configured to receive, organize, and store in the non-volatile memory information that enables execution of at least one service procedure for the industrial machine. In example embodiments, the system above may also include a data storage element accessible through a processor, the data storage element comprising a copy of information stored in a plurality of the smart RFID element. In example embodiments, each copy of information comprises a twin of the information stored in the corresponding smart RFID.

[1810] In example embodiments, industrial machine predictive maintenance methods and systems, such as those described herein may include use of one or more machine-resident smart RFID data structures that may capture information related to planning, engineering, production, assembly, testing and the like of portions of the industrial machine. Embodiments 33100 that may facilitate capturing information from these processes may be depicted in Figure 185. An industrial machine 33122 may comprise several elements, such as operational elements, structural elements, processing elements, and at least one smart RFID element 33102. During production of the industrial machine 33122, an industrial machine-resident processor 33108 may work cooperatively with self-test elements 33124 and the like to perform testing of the industrial machine. Data collected during self-testing, such as confirmation of proper operation and the like may be stored in the smart RFID element 33102, such as by the processor writing this data into a memory of the smart RFID element 33102. In example embodiments, a production test system 33118 may also perform testing of portions of the industrial machine 33122, the results of which may be stored on the smart RFID element 33102. The industrial machine 33122 may communicate with a production network 33120, such as an intranet and the like during production to gather and/or provide information for various production systems, such as quality systems 33110, manufacturing resource and planning (MRP) systems 33114, production engineering systems 33116 and the like. Information, such as parts lists, production information, and the like, an example data structure of which is depicted in Figure 182, may be stored with the smart RFID element 33102, such as by the industrial machine 33122 communicating over the production network 33120 via a production access point 33112 and the like. Information from the various production systems, quality 33110,

MRP 33114, engineering system 33116, testing 33118 and the like may be transferred over the network 33120 to the smart RFID element 33102. In example embodiments, a networked server 33126 may communicate with at least a portion of these production systems over the network 33120 to, for example capture and process with a processor 33106 relevant production information to be stored in the smart RFID element 33102 and/or in a data structure in a memory accessible to the server 33126. A data structure 33104 may include at least a portion of the information stored in the smart RFID element 33102. In example embodiments, the data structure 33104 may be a digital twin of at least the relevant production content of the smart RFID element 33102 for the specific industrial machine being produced. In example embodiments, data from the production systems may flow through the network 33120 to the server 33126 and may optionally be processed there, such as to be formatted, encoded, and the like and delivered, such as over a wireless connection to the industrial machine 33122 for storing with the smart RFID 33102. Production systems may include the quality control systems 33110 that may include capturing images of parts, sub-assemblies, and portions of the industrial machine. Images captured may be processed with machine vision and other image analysis technologies to validate assembly and the like. These images, image analysis data derived from these images, and the like may be stored so that it may be accessed through the smart RFID element 33102. In an example, procedures such as test procedures used in production may be useful for testing the industrial machine 33122 as part of a deployment process. These procedures may be communicated from one of the production systems, such as the engineering system 33116 over the production network 33120, eventually to be stored on the smart RFID 33102, the digital twin 33104 or both. This may satisfy a goal of the methods and systems described herein of facilitating access to industrial machine-specific procedures via a smart RFID element on each industrial machine.

[1811] In example embodiments, production information stored in, for example the smart RFID element 33102 may be useful to procedures that are to be followed during installation, calibration, repair, preventive maintenance and the like. In an example, certain test results may indicate an operational margin (e.g., maximum and/or minimum values) verified during production. These results may be useful during validating testing of a deployment of the industrial machine to facilitate confirming the deployment continues to meet expectations. By making this and other production and industrial machine information available during installation and other deployed procedures, the machine-resident smart RFID element 33102 reduces interdependency of production and related systems once an industrial machine leaves the production environment. In an example, a procedure for testing a portion of the industrial machine may be stored in the smart RFID element. Test results that correspond to that procedure may also be stored therein. Therefore, even if the specific procedure is modified for subsequently produced industrial machines, it may be possible to perform tests associated with the specific procedure used to produce the specific industrial machine; thereby saving time and confusion that may occur when a new test procedure is used, but old procedure test results are expected to be met.

[1812] In example embodiments, a method of configuring production data in a smart RFID of an industrial machine may include configuring a smart RFID with a portion of an industrial machine

to capture and store in a non-volatile computer-accessible memory operational, physical and diagnostic result information for a corresponding portion of the industrial machine. The method may include communicatively coupling the smart RFID with a processor of the industrial machine and at least one sensor configured to monitor a condition of the portion of the industrial machine. The method may further include executing with the processor a self-test of the portion of the industrial machine and storing in the smart RFID a result of the self-test. The method may yet further include coupling the industrial machine through a production access point to a network of testing systems and an industrial machine production server. The method may further include performing production tests on the portion of the industrial machine with the testing systems, a result of which is stored in duplicate on the smart RFID and in a data storage facility accessible by a processor of the production server. In example embodiments, the duplicate of the testing results stored in the data storage facility may be a twin of the corresponding portion of the smart RFID.

[1813] In example embodiments, a marketplace of industrial machine parts, services, tools, materials and the like may be maintained through a combination of a CMMS control system, and third parties each providing information about services, parts, tools, materials, costs, and logistics that they provide. Such a marketplace may be cloud-based so that access to this information, can be made available to participants including industrial machine owners and the like. In example embodiments, a representative embodiment is depicted in Figure 186. A CMMS system 33202 for managing at least part and service orders for required services may act as a control gateway to a marketplace 33212 for industrial machine owners 33224 and the like. The CMMS system 33202 may include managing bids and orders for parts, service, tools, materials and other aspects of industrial machine service and maintenance. Exemplary CMMS subsystems, systems, facilities and the like are described elsewhere herein. In the embodiment of Figure 186, the CMMS system 33202 may further maintain and update order history details 33210. These details may include information descriptive of the parts, services, and the like that may be ordered. Details may include historical pricing, logistics requirements and costs, order lead times, and other factors that may be useful when managing information in the marketplace 33212. In an example, a part supplier 33208 may offer a part for sale in the marketplace. Historical pricing for the part based on the order details 33210 may be used to recommend a price at which the part supplier 33208 should offer the part. In another example, the part supplier 33208 may offer availability of a part with a 2-day lead time. However, the historical details 33210 may indicate that this supplier 33208 is underestimating the time required to provide the part and may facilitate incorporating a proper lead time when placing the order so that the part can be ordered only when needed but with sufficient lead time for it to be available when a service that requires the part is scheduled to be performed. Such information management may be implicit management because it is based on actual performance rather than mere statements by a provider.

[1814] In example embodiments, service providers 33206 may configure offering for a set of services 33216 that meet their technical expertise. The service providers 33206 may directly configure and update this set of services over time so that it reflects the services available from each individual service provider 33206 over time. Likewise, the parts supplier 33208 may

configure and maintain a list of parts 33214 for industrial machines that the supplier offers. Information such as availability (e.g., local inventory, lead time, and the like) may be directly maintained by the parts supplier 33208. The CMMS system 33202 may access his and related information in the marketplace 33212 when configuring an order for parts, services, and the like. Similarly, suppliers of tools may configure information regarding industrial machine service tools 33220 and suppliers of materials may configure and maintain information regarding industrial machine service materials 33222 (e.g., lubricants, other consumable items, and the like).

[1815] In example embodiments, parts manufacturers 33204 may also provide and maintain information regarding parts that they provide, such as replacement parts, add-ons, upgrades, complete systems, subsystems, accessories and the like to the marketplace.

[1816] In example embodiments, a logistics suppliers 33218, such as shippers and the like, may provide and maintain a set of logistics services in the marketplace that they provide for industrial machine maintenance parts, services and the like. The logistics supplier 33218 may offer delivery services in different geographic regions and may use information such as location of the industrial machine to establish rates and services available in the relevant region.

[1817] In example embodiments, an industrial machine predictive maintenance system may form a marketplace that includes a plurality of parts supplier computing systems configured to maintain industrial machine service marketplace information about industrial machine parts offered for sale. The marketplace may include a plurality of service provider computing systems configured to maintain industrial machine service marketplace information about industrial machine services offered. The marketplace may further include at least one computerized maintenance management system (CMMS) that is configured to facilitate access to at least one of services, parts, materials, and tools offered in the marketplace responsive to an industrial machine maintenance recommendation provided by an industrial machine predictive maintenance system. The marketplace may yet further include a plurality of logistics provider computing systems configured to maintain industrial machine service marketplace information for at least one of shipping and logistics services offered in the marketplace. Further In example embodiments, each of the plurality of parts suppliers, service providers, and logistics providers maintain corresponding information for their offerings directly in the marketplace via at least one Application Programming Interface of the marketplace. The marketplace may further include a CMMS that adapts offerings of parts, services, and logistics to industrial machine owners based on norms established from analysis of prior orders for parts, services and logistics.

[1818] In example embodiments, a distributed ledger for tracking field service activities, including predicative maintenance activities and the like that are performed on industrial machines is depicted in FIG 187. Methods and systems that are disclosed herein for an industrial machine maintenance distributed ledger may include a distributed ledger 33302 supporting the tracking of predictive maintenance activities executed in an automated industrial machine predictive maintenance eco-system 33300. Embodiments may include a self-organizing data collector 33308 that is configured to distribute collected information to the distributed ledger 33302. Embodiments may include a network-sensitive data collector that is configured to distribute collected information

to a distributed ledger based on network conditions. Embodiments may include a remotely organized data collector that is configured to distribute collected information to a distributed ledger based on intelligent, remote management of the distribution. Embodiments may include a data collector with self-organizing local storage that is configured to distribute collected information to a distributed ledger. Embodiments may include the system 33300 for industrial machine maintenance-related data collection in an industrial environment using a distributed ledger for data storage and self-organizing network coding for data transport. In example embodiments, data storage may be of a data structure that supports a haptic interface for data presentation, a heat map interface for data presentation, and/or an interface that operates with self-organized tuning of an interface layer.

[1819] In example embodiments, storage of service and maintenance information, which may include services, parts, service providers, records for specific industrial machines, analytics generated from the service and maintenance information and the like may include the one or distribute ledger 33302 instances in various elements of the system 33300. In an example, the distributed ledger 33302 may facilitate access to all of the information available in the distributed ledger 33302 without relying on any one network server, node, or the like due at least in part to some portion of the information being distributed and optionally duplicated on distinct portions of a network, such as the Internet. The distributed ledger 33302 may be distributed among elements in an industrial machine maintenance platform including, without limitation, the industrial machine data analysis system 28602, the industrial machine predictive maintenance subsystem 28616, the CMMS system 28622, the service delivery and tracking system 28630, the industrial machine 33304, the industrial facility computing system 33306, the cloud-based storage 33316, and the like.

[1820] In example embodiments, information stored in the distributed ledger 33302 may be generated by and/or adjusted based on artificial intelligence 33310, such as machine learning algorithms that process the information from which the distributed ledger is sourced.

[1821] In example embodiments, the methods and systems that may support distributed ledger embodiments may include role-based access control 33314 of and to the distributed ledger data. Exemplary roles 33312 that may be managed by a distributed ledger control facility may include: an owner role, which may be an industrial machine leasing company, individual or direct-use buyer entity or individual; an operator role, which may be an entity or individual that is responsible for day to day operation of an industrial machine, such as a company that provides a service using the industrial machine, a lessor of the machine, and the like; a lessor role, which may be an entity or individual that has a term-based or otherwise limited lease of an industrial machine; a manufacturer role, which may be an entity or individual that produced some portion of the machine and that may have limited access to, for example, information pertaining to the portion produced; a part supplier role, which may be an entity or individual that provides some part(s) for manufacturer, service, upgrade, maintenance, refurbishing, or other functions and may provide OEM and/or after-market parts for an industrial machine; a service provider, which may be an individual or entity that provides services, such as contracts for preventive maintenance and repair, emergency repair, upgrades and the like; a service broker role, which may be an entity or individual that facilitates

service needs, such as a regional entity that facilitates automated service activities in regions, such as specific countries and that may be required to be licensed, registered, and the like in the specific country and that may act comparably to a general contractor, providing oversight and warranty for work done by 3rd parties, such a role may be valuable when a machine has been installed per local rules, and the like that is outside of the scope of what an automated service identification system may handle; a regulatory role, which maybe a government or other authority entity or individual that may conduct inspections and the like and may be limited to access certain data required for ensuring compliance with regulations and the like for activities such as preventive maintenance, use of authorized parts/service providers, auditing, and the like.

[1822] In example embodiments, a predictive maintenance platform may use a secure architecture for tracking and resolving transactions, such as a distributed ledger. In example embodiments, transactions in data packages are tracked in a chained, distributed data structure, such as a Blockchain™, allowing forensic analysis and validation where individual devices store a portion of the ledger representing transactions in data packages. The distributed ledger may be distributed to IoT devices, to web servers, to industrial machine maintenance transaction record storage facilities, and the like, so that maintenance and related information can be verified without reliance on a single, central repository of information. The platform may be configured to store data in the distributed ledger and to retrieve data from it (and from constituent devices) in order to resolve service transactions, such as parts and service orders, and the like. Thus, a distributed ledger for handling data for maintenance-related transactions is provided. In example embodiments, a self-organizing storage system may be used for optimizing storage of distributed ledger data, as well as for organizing storage of packages of data, such as IoT data, industrial machine maintenance data, parts and service data, knowledgeable worker data, and the like.

[1823] In example embodiments, a system may include a plurality of computing systems configured to perform one or more predictive maintenance actions. In example embodiments, a portion of the plurality of computing systems connected via a peer-to-peer communication network. A record of industrial machine maintenance actions including a portion of the predictive maintenance actions may be maintained by the portion of the plurality of computing systems as a distributed ledger. In example embodiments, a computing system of the portion of computing systems performs at least one industrial machine maintenance role selected from the list consisting of industrial machine data analysis, industrial machine predictive maintenance recommendations, industrial machine maintenance order management, delivery and tracking of service actions, industrial machine service scheduling, and contributes a result of it performing the at least one industrial machine maintenance to the record.

[1824] In example embodiments, a system may include a plurality of computing systems configured to perform one or more predictive maintenance actions. In example embodiments, a portion of the plurality of computing systems are connected via a peer-to-peer communication network. In example embodiments, the system may further include a role-based control facility for accessing a record of industrial machine maintenance actions, the record including a portion of the

predictive maintenance actions. In example embodiments, the portion of the plurality of computing systems operate the record as a distributed ledger.

[1825] In example embodiments, methods and systems for operating a predictive maintenance analysis and control system may benefit from visual information as well as performance and operational data from industrial sensors and the like deployed with an industrial machine. Visual information, such as images captured about individual parts, assemblies, process steps, machine conditions and the like may be analyzed with machine vision and other techniques, including human viewing and assessment, to determine conditions that may impact prediction of a service need or the like. Generating and maintaining an updated accurate image library of visual information for industrial machines may be benefited from service personnel capturing images of portions of each industrial machine under various conditions, including without limitation operating, testing, and non-operating conditions (e.g., during service, maintenance, repair, upgrade, and refurbishing machine states). In example embodiments, a system to facilitate capture of images is depicted in Figure 188. A procedure for industrial machine service or repair 33416 may be identified for a scheduled service of the machine. The procedure 33416 may include a set of steps to be taken to perform the scheduled service activity. One or more of the steps may include capturing image(s) of portions of the industrial machine, such as an external view depicting the machine in its deployed environment, a view of a part to be replaced, a view depicting a condition of gears, bearings, support structures, housings and the like. While a procedure may include capturing image(s), learning from service technicians performing the procedure may be incorporated into implementing the procedure using a preventive maintenance system 33424 that uses machine learning and other techniques to facilitate augmenting and/or adjusting image capture steps in a procedure and the like. The predictive maintenance system 33424 may provide information, such as in the form of conditions that suggest an image should be captured that may not be directly required in a procedure. Such a case may arise when the predictive maintenance system 33424 learns that certain bearings exhibit wear that is visible before the bearing fails. The length of time that a bearing can operate under various conditions may not be a sufficient indicator to perform a service, whereas an image with visual indication of such wear would be sufficient. Therefore, when a service technician performs a service procedure that does not include capturing an image of the certain bearings, the technician may be directed to capture an image of these certain bearings. This may be indicated to the service technician as a service alert, such as a general posting. However, information about the visual condition and timing of a service activity may be used to facilitate augmenting / updating a procedure, such as the procedure 33416 to include capturing one or more images of the certain bearings.

[1826] In example embodiments, information from the predictive maintenance system 33424 may be processed by an image capture triggering facility 33422 to provide an indication to a procedure updating facility 33402 that an update to the procedure, such as to add capturing an image of the certain bearings, is required. This indication may be combined with image capture timing information that may be provided to the procedure update facility 33402 from an image capture timing facility 33420 that may use industrial machine use and service schedule information 33426

to create a window of time in which the certain bearings are expected to be available to be imaged. Such a window of time may include scheduled service and/or maintenance activities during which the machine may be off-line. Such a window of time may include planned operational times during which the machine will be operating. A potential goal of such window generation may be to capture image(s) of the certain bearings during a planned service visit, to avoid machine shut downs specifically to capture the image(s), despite the images being required before a service activity in which the bearings would normally be images is executed, such as a scheduled preventive maintenance activity to inspect the bearings and the like.

[1827] In example embodiments, when the existing procedure 33416 is to be applied during an image capture window output from the image capture timing facility 33420, the image capture triggering facility 33422 output may be checked. If the image capture triggering facility 33422 indicates that an image is required, the procedure may be updated by the procedure update facility 33402, such as by adding a step to the procedure, changing an imaging target (e.g., from a part to the bearings) for an existing image capture step, and the like.

[1828] In example embodiments, the revised procedure 33402 may be followed by the service technician. When a step that has been added/augmented to capture an image of the certain bearings is to be performed, an image capture template 33404 may be presented to the technician to aid in capturing the proper image. Likewise, and as described elsewhere herein, an augmented reality application may be executed as part of such an image capture step to further aid the service technician in capturing the proper image. In example embodiments, a machine vision system 33408 and other image analysis techniques may be used to suggest refinements and/or confirm the captured image meets the requirements for facilitating detecting the visual condition of the certain bearings, and the like.

[1829] In example embodiments, an image capture reward facility 33414 may interface with the updated procedure 33418 and/or the service technician to facilitate incentivizing the service technician to capture an acceptable image. Such a reward facility 33414 may include a range of rewards from direct monetary rewards to positive ratings for the service technician, which may ultimately increase the technician's value and consequently compensation.

[1830] Captured images, such as those that are accepted by the machine vision system 33408 and the like, may be stored in a smart RFID element 33410 of the industrial machine, transferred through the image capture device (e.g., a camera-enabled smart phone, and the like) to the Smart RFID and to one or more nodes in a distributed ledger of preventive maintenance data.

[1831] In example embodiments, a method of image capture of a portion of an industrial machine includes updating a procedure for performing a service that implements a predicted maintenance action on an industrial machine, the updating responsive to a trigger condition for capturing an image of a portion of the industrial machine being met. The method of image capture may further include providing an image capture template in an electronic display overlaying a live image of a portion of the industrial machine to facilitate image capture, applying augmented reality that indicates a degree of alignment of the live image with the template, examining an image captured using the updated procedure with machine vision to determine at least one part of the machine

present in the captured image, and responsive to a result of the machine vision examination, operating an image capture reward facility to generate a reward for the captured image. In example embodiments, the updating may be responsive to a trigger condition that is based on analysis of industrial machine failure data such that the analysis suggests capturing an image that is not specified in the procedure prior to the updating step. In example embodiments, the updating may be responsive to the procedure for performing the service being performed on an industrial machine that meets a predictive maintenance criterion associated with the portion of the industrial machine for which an image is to be captured. In example embodiments, the trigger condition may include a type of industrial machine associated with the industrial machine for which a service procedure is being performed and a duration of time since the portion of the industrial was captured in an image.

[1832] In example embodiments, an industrial machine predictive maintenance facilitating system may apply machine learning to images of industrial machines captured during operations such as assembly, testing, servicing, repair, upgrading, scheduled maintenance, preventive maintenance, and the like. The machine learning may be applied to the images and/or data derived from the images using algorithms such as image analysis algorithms, part detection algorithms, machine vision and the like to facilitate improving machine-automated detection of portions of the industrial machine, such as individual parts, subassemblies and the like. In example embodiments, machine-automated detection of parts, subassemblies and the like may provide information to the methods and systems here including, without limitation, predictive maintenance processes, service provider rating methods, procedure rating methods, inventory management systems, maintenance scheduling (e.g., if a maintenance operation should be scheduled sooner than previously estimated, and the like).

[1833] In example embodiments, methods and systems for machine-automated detection of parts of an industrial machine may include image capture, processing, analysis, learning and automation steps, such as those exemplarily depicted in Figure 189. In example embodiments, a method for automatically detecting parts of an industrial machine may start with capturing an image step 33502. Alternatively, image data from previously captured images may be accessed from a data store of images, such as a database and the like. The image capture step 33502 may be performed, such as by a service technician and the like in association with performing a service operation, such as a maintenance procedure, repair procedure, upgrade procedure and the like. The image capture step 33502 may be informed by a procedure or the like that may indicate a target part to be imaged, a template thereof, and the like. A procedure, target part, template and the like may be retrieved from an image capture guidance data storage 33504. In example embodiments, a procedure may include a specific instruction to use a part image capture process and photograph one or more parts indicated by the procedure. In an example, a procedure for servicing bearings of an industrial machine may include a step of photographing a shaft that the bearings handle and the like. The procedure may be present on an electronic display of an image capture device, such as a tablet or smart phone and the like an image representative of the image to be captured. Such an image may be a most recent image captured of the specific industrial machine that may, for

example, be retrieved from an image data structure of a smart RFID element deployed with the industrial machine (e.g., a smart RFID element configured with the portion of the machine that includes the bearings, shaft and the like). Such an image may be augmented with information, such as relative position of the camera through which the image was captured, time/date information, procedure number followed, and the like. In example embodiments, such an image may be processed into a template (e.g., coloring book / outline image, and the like) that facilitates manually aligning the image capture device. In example embodiments, such a template may be an active template that processes an image visible through the image capture device and provides indicators, such as color changes and the like of the template to further facilitate alignment of the image capture device. The active template may start with black (or some other color) outlines of the object(s) to be captured with vertexes, edges, and the like turning green (or some different color) when alignment of the relevant vertex, edge and the like is sufficient to facilitate machine-automated detection of the part.

[1834] In example embodiments, an image captured in the image capture step 33502 may be processed through an image validation step 33506 that may perform image analysis functions, such as for example comparing the image captures with a reference image, such as one that may be retrieved from or derived from information in the image capture guidance data store 33504 and the like. In example embodiments, the captured image may be processed to improve contrast and the like and compared during the validate image capture step 33506 with a most recently captured image from the smart RFID element disposed with the industrial machine through, for example an image subtraction process, to determine if the captured image may be validated. An image that is not validated may be discarded and the user may be directed back to the capture image step 33502 to capture another image.

[1835] In example embodiments, an image that may be validated in step 33506 may be passed onto an image analysis or a similar step 33508 that may process image analysis rules 33510 to detect one or more candidate parts from the validated image. Candidate parts may be stored in a candidate parts data structure 33514 for further use. In example embodiments, images of candidate parts in the candidate parts data structure 33514 may be retained for further training of machine learning algorithms that facilitate improving machine automated part detection from images. In example embodiments, images of candidate parts may be used in an instance of the machine automated parts detection flow 33500 of Figure 189 and then discarded, erased, and the like. In example embodiments, the image analysis rules 33510 may include data provided from the machine learning step 33520, such as in the form of feedback and the like that may improve image analysis of marginal images, such as those with poor contrast, unexpected content (e.g., excessive solvents, moving parts, reflective parts, and the like).

[1836] In example embodiments, the one or more candidate parts of the candidate parts data structure 33514 may be processed by a parts recognition algorithm step 33516 that may perform, among other things, machine automated parts recognition. An automated parts recognition algorithm may include generating attributes of candidate parts, such as dimensions and the like that may be compared with part descriptive information that may be retrieved from a smart RFID

data storage 33512, and the like. In an example, a candidate part may be processed to detect edges and the like that may be processed with automated measurement algorithms. The resulting measurements may be used to determine a specific part from a library of parts for the specific industrial machine that may be available to the parts recognition algorithm 33516 in the RFID data storage 33512 and the like. The specific part information may be retrieved from a production data system, such as a parts list, MRP system and the like and stored in the RFID data storage 33512 during a production operation, such as the exemplary production flow depicted in Figure 185.

[1837] In example embodiments, one or more results of the parts recognition algorithm 33516 may be forwarded to a machine learning facility, that may execute one or more machine learning algorithms 33520 that may improve various aspects of machine-automated part detection including, without limitation, the image capture process 33502, the image validation process 33506, the image analysis process 33508, the part recognition process 33516 and the like. In an example, part recognition process 33516 may provide images of one or more candidate parts, a corresponding reference part, related attributes and the like, information extracted during the parts recognition process, and the like to the machine learning process 33520. The machine learning process may apply machine learning techniques to facilitate determining aspects of candidate part(s) that represent the best candidates for the corresponding reference part and provide feedback to at least the part recognition process 33516 to improve part detection and the like.

[1838] In example embodiments, information descriptive of recognized parts may be stored in an updated smart RFID element 33518, an updated server-based data structure 33522 comparable thereto, and the like. Information stored may include one or more candidate part images, an identifier of a reference part, recognition data, procedure number followed to capture the image, and the like.

[1839] In example embodiments, a method of machine learning-based part recognition may include applying a target part imaging template to an image validating procedure that determines if an image captured meets an image capture validation criterion. The method may further include performing image analysis by processing a captured image with image analysis rules that facilitate detecting candidate parts of an industrial machine being present in an image. In example embodiments, recognizing one or more parts of the set of candidate parts as a part of the industrial machine based on similarity of a candidate part with images of parts of the specific industrial machine may be included. Additionally, adapting at least one of the target part template, the image analysis rules, and the part recognition based on feedback produced from machine learning of the recognized parts, thereby improving at least one of image capture, image analysis and part recognition may be included in the method.

[1840] In example embodiments, information gathered and generated for industrial machine maintenance lifecycles, including predictive maintenance, manufacturer required maintenance, failure repairs, parts and service offerings and ordering, follow-up to maintenance activities, assessment of procedures and service providers, failure rate and prediction analysis, worker training, experience, and ratings, and the like may be captured throughout the service lifecycle, processed with artificial intelligence and other machine learning-type algorithms and accumulated

in a database, such as a data model, linked database, columnar database, and the like. Figure 169 depicts such a set of data embodied as a knowledge graph 33602. In example embodiments, information about industrial machines, such as parts, images, configurations, internal structures, use schedules, and the like may be processed by artificial intelligence-type functions 33606 (e.g., machine learning algorithms and the like) along with information from other sources including without limitation service information, failure information, worker-related information and the like. The information processing algorithms, such as information associative algorithms executed in exemplary artificial intelligence facility 33606 may cause portions of the predictive maintenance and industrial machine service knowledge graph 33602 to be updated, such as by establishing, changing, removing, strengthening and the like knowledge graph node links 33616 among data nodes 33618; adding, updating, splitting and the like the data nodes 33618 to initiate and refine a graph-based understanding of the relationships among facts, know-how, analysis results and the like that influence aspects of predictive maintenance processes, such as those described herein.

[1841] In example embodiments, information about machines may be processed and stored in machine data nodes 33608; information about failures may be processed and stored in failure data nodes 33610; information about industrial machine service may be processed and stored in service data nodes 33612, information about workers for performing industrial machine service may be processed and stored in worker data nodes 33614. Relationships among data nodes, such as a relationship between the machine data node 33608 and the service data node 33612 may be depicted as the links 33616 between nodes. A goal of initiating and updating such a knowledge graph, among other things may be to further improve for collecting, discovering, capturing, disseminating, managing, and processing information about industrial machines, including factual information (such as about internal structures, parts and components), operational information and procedural information, including know-how and other information relevant to maintenance, service and repairs.

[1842] In example embodiments, as maintenance/service/repair/upgrade/installation and other industrial machine-related activities are performed, data about the activities may be processed and used to enhance, augment, improve, refine, clarify, and correct the data nodes 33618, the relationships among the nodes, and the like. In example embodiments, preparing for maintenance/service/repair and other industrial machine activities may benefit from the knowledge found in the knowledge graph 33602 and thereby improve efficiency, reduce computing complexity to generate suitable service options, recommendations, orders and the like by taking, for example an existing relationship between the failure node 33610 and the worker node 33614 to efficiently identify a suitable worker for resolving the failure when it occurs on a specific machine.

[1843] In example embodiments, improved methods and systems are provided herein for collecting, discovering, capturing, disseminating, managing, and processing information about industrial machines, including factual information (such as about internal structures, parts and components), operational information and procedural information, including know-how and other information relevant to maintenance, service and repairs. These improved methods and systems may be provided with a predictive maintenance knowledge system platform 33700 as depicted in

Figure 191. A predictive maintenance knowledge system 33702 may facilitate collecting, discovering, capturing, disseminating, managing, and processing information about industrial machines, such as for facilitating service and maintenance thereof using the methods and systems described herein, including without limitation finding a set of workers having relevant know-how and expertise about maintenance, service and repair of a particular machine and finding, ordering, and fulfilling orders for relevant parts and components, so that maintenance, service and repair operations can occur seamlessly, with minimal disruption, and the like. The predictive maintenance knowledge system 33702 may interface with one or more predictive maintenance knowledge bases and/or knowledge graphs 33704. A knowledge base 33704 may further include or reference one or more knowledge graphs that may contain information beneficial to the methods and systems that may be enabled by the predictive maintenance knowledge system 33702. The predictive maintenance knowledge graph may contain or provide computer access to information about industrial machines, service activity of industrial machines, costs (e.g., historical, trending, and predictive) for parts, materials, tools, and services of industrial machines, algorithms and functionality for operating the predictive maintenance knowledge system 33702, platform 33700 and the like. In example embodiments, the predictive maintenance knowledge system 33702 may process information from the predictive maintenance knowledge base 33704 regarding expedited service charges that have been imposed on certain instances of industrial machine service and develop a price-time relationship that may aid in the decision by industrial machine owners regarding service authorization and costs thereof. An industrial machine owner may be informed of the costs for expedited service and standard timing service to facilitate deciding if it is better to pay an expedite fee to have a maintenance function performed soon while the machine is off-line for other reasons than to keep a schedule of the maintenance function that would require taking the machine off-line, such as in the near future. The predictive maintenance knowledge system 33702 may facilitate coordination with service providers, parts providers, material and tool providers and the like based on the owner's decision so that the service can be performed in the timeframe that the owner chooses.

[1844] In example embodiments, specific industrial machine information may be stored in one or more smart RFID elements 33706 disposed with the specific machine and/or stored in a cloud-based data structure 33708 that may be compatible with (e.g., a backup, duplicate/twin, or other formatted data structure). The predictive maintenance knowledge system 33702 may access (e.g., read data from and/or write data to) the RFID element(s) 33706, the cloud-based data structure 33708, and the like. Data read from the smart RFID 33706 / cloud-based structure 33708 may be specific to a particular deployed industrial machine and may facilitate the methods and systems for predictive maintenance and the like described herein performing coordination of resources to perform maintenance effectively and efficiently for the specific machine. In an example, a specific industrial machine may have an operating cycle that results in greater utilization of one of its moving parts (e.g., an industrial motor) than typical. This knowledge may be used by the predictive maintenance knowledge system 33702 to interact with service, parts, and material suppliers to provide a firm quote for performing a utilization-based maintenance service at a different time

(e.g., weeks or months sooner) than other comparable industrial machines with lower utilization rates.

[1845] In example embodiments, the predictive maintenance knowledge system 33702 may execute algorithms that gather information about a plurality of industrial machines, including a plurality of industrial machines of different types of machine (e.g., stationary machines, mobile machines, machines on vehicles, machines deployed at job sites, and the like) along with service provider information, parts and parts provider information, part location and inventory information, machine production providers, third-party parts handlers, logistics providers, transportation providers, service standards, service requirements, service activities including results of service and the like, and other information to facilitate the predictive maintenance methods and systems described herein. One or more functions of the predictive maintenance knowledge system 33702 may utilize service request information 33726, such as requests for service of a specific industrial machine and/or a collection of industrial machines from industrial machine owners/operators/providers/users to facilitate fulfilling those service requests. In example embodiments, such service requests may become inputs to an algorithm that predicts when a service may be recommended for the requester, but also for comparable industrial machines. In an example, an industrial machine owner may request that a subset of industrial machines at a job site receive a first service action. The predictive maintenance knowledge system 33702 may use this request information and other information about the machines, such as their age and utilization rate, to determine when the other industrial machines of the same type as those for which the service is requested should be scheduled for a comparable service action.

[1846] In example embodiments, in response to the specific service request 33726, the predictive maintenance knowledge system 33702 may access information in the smart RFID 33706 or its cloud-based backup 33708 to determine the specific procedures involved, to determine what experience a potential service provide may need to perform the service. The predictive maintenance knowledge system 33702 may access the knowledge base 33704 to identify candidate service providers. Service providers that are known to the predictive maintenance knowledge system 33702 (e.g., based on, for example information in the knowledge base 33704) as having successfully demonstrated experience with the procedure needed for the requested service may be contacted to provide a service estimate 33736 and/or a price estimate 33734 for service, parts, and the like. Similarly, parts and/or material that may be associated with the procedure of the requested service may be identified. The predictive maintenance knowledge system 33702 may also access the knowledge base 33704 for sourcing information of the parts and/or material. Factors such as part cost, transportation costs, availability, location of the parts versus the machines, prior relationships between one or more parts providers and a party associated with the service request, such as the industrial machine owner and the like, and other factors may be evaluated to determine which parts provider to contact in preparation for ordering the parts. With these factors considered, a part inquiry may be placed with one or more parts providers in anticipation of the service being conducted by the qualified service provider as scheduled. The predictive maintenance knowledge system 33702 may respond to the service request 33726 with one or more service recommendations

33732 that may be associated with one or more price-based service recommendation options 33710 from which the requestor may choose. In example embodiments, the predictive maintenance knowledge system 33702 may have enough information from the knowledge base 33704, responses to the service estimate request 33736, and the like to automatically select a specific price-based service recommendation 33710 from the options and may, with or without requestor explicit approval, generate a service order 33718, a parts/material/tools order 33716 if needed for the requested service 33726.

[1847] In example embodiments, a service request and/or a predicted maintenance activity, and the like may be processed by the predictive maintenance knowledge system 33702 and output a service funding recommendation and/or request 33712. Such a recommendation may include funding the service from operating revenues, taking out a loan for the service, seeking third-party funding (e.g., industry sources, government grants, private funding sources, and the like). Such a request may include providing information to one or more third-parties about the requested service that may be used by the third-parties to submit a funding proposal and/or response. In an example, an industrial machine that provides the public with clean water for a region may require a costly service. The predictive maintenance knowledge system 33702 may determine that the specific industrial machine may be eligible for reimbursement from the federal government for at least a portion of the service. A request for funding by the federal government may be configured and activated through the service funding 33712 and the like.

[1848] In example embodiments, sources of information that the predictive maintenance knowledge system 33702 may rely on may include information from service providers 33724, information from parts providers 33722, information from service material providers 33720, machine schedules 33730, incoming service estimates and/or quotes 33728, and the like. A predictive maintenance knowledge system 33702 may use service material provider information 33720 to determine price and availability of service material. This information may be combined with service material inventories of the requester (e.g., centralized, depot-based, or on-site of the industrial machine), inventories of material of one or more qualified service providers and the like. In an example, if a service provider has sufficient inventory of the required material accessible local to the industrial machine for which service is required, but will need to replenish that inventory after performing the service, the system may provide a recommendation to the service provider to have the service material provider deliver the service material to the industrial machine site in time for the schedule service. In an example, if the service provider and the industrial machine owner does not have inventory of the required service material, the predictive maintenance knowledge system 33702 may generate an order with one of the service material providers 33720 based on total price, availability, existing relationships with the industrial machine owner and/or the service provider and the like. In example embodiments, at least a portion of the inventory of one or more of the service material providers 33720 may be directly managed by the predictive maintenance knowledge system 33702 so that the predictive maintenance knowledge system 33702 may allocate material from the inventory for a service action. The service material provider 33720 may receive a notification from the predictive maintenance knowledge system

33702 that they have been selected to provide the material for the service action. Payment for the material may be made through a transaction facility associated with the predictive maintenance knowledge system 33702 so that an operator of the predictive maintenance knowledge system 33702 and the service material provider 33720 are compensated for their roles in this service action. Comparable examples may be envisioned for parts providers 33722, service provider 33724, service funding sources (not shown), and the like.

[1849] In example embodiments, the predictive maintenance knowledge system platform 33700 may include a computerized maintenance management system (CMMS) 33714 that may facilitate creating work orders, such as for maintenance actions to resolve equipment problems, and the like. The CMMS 33714 may facilitate communicating parts and service requests to an Enterprise Resource Planning (ERP) system (not shown) that may facilitate handling parts and service orders. In example embodiments, an ERP system may be associated with one or more of the owner/operator/provider/lessee/lessor of an industrial machine for which a service action is being coordinated by the predictive maintenance knowledge system 33702. In example embodiments, the CMMS 33714 may coordinate with the industrial machine owner's ERP system to effect placement of orders with the service provider, parts provider, and the like.

[1850] In example embodiments, a predictive maintenance system may include a predictive maintenance knowledge system that facilitates collecting, discovering, capturing, disseminating, managing and processing information about industrial machines to facilitate taking predictive maintenance actions on industrial machines. The knowledge system may include a plurality of interfaces for receiving information from service providers, parts providers, material providers, machine use schedulers, a plurality of interfaces for sending information to service ordering facilities, parts ordering facilities, service management facilities, service funding facilities, and a plurality of interfaces to smart RFID elements on a plurality of industrial machines. The predictive maintenance system may further include a predictive maintenance knowledge graph that facilitates access by the predictive maintenance knowledge system to information about predictive maintenance service of industrial machines through links among data domains including service providers, parts providers, service requests, service estimates, machine schedules, and predictions of maintenance activity. In example embodiments, the predictive maintenance knowledge system may generate at least one of service recommendations, price-based service options, price estimates, and service estimates.

[1851] In example embodiments, preventive maintenance and other scheduled maintenance for industrial machines and the like may be scheduled at set intervals based on manufacturer's expectations regarding failure rates and the like. By gathering and analyzing information about industrial machines and the like, such as operational data, failure data, conditions found during preventive maintenance activities and the like, a new schedule for maintenance activities may be configured that may further reduce the occurrence of unplanned shutdowns due to part failure and the like. Figure 192 depicts a preventive maintenance schedule 33808 for a set of bearings in a group of industrial machines 33802 that use the bearings. As presented, preventive maintenance events A, B, C, and D for the bearings are scheduled to occur at intervals over time for each of the

machines. Data collected and analyzed by a predictive maintenance system using the methods and systems for predictive maintenance of industrial machines as described herein may indicate that a different schedule of bearing maintenance is needed to prevent failures. In the example of FIG 192, failures 33804 of machines 4 and 3 occur after preventive maintenance activity B. In response there to, and when taking into consideration other factors, such as operating cycle rate of the industrial machines, a new bearing maintenance schedule may be established for the machines. Since machines 1 and 2 have not yet failed, a predictive maintenance event may be set up for machine 1 33810 and for machine 2 33812. In example embodiments, an operational rate of machine 2 may be substantive less than machine 1; therefore, while both machines use the bearings that have failed in machines 3 and 4, a predictive maintenance event schedule may be prepared individually for each machine. The predictive maintenance event for machine 1 33810 may be set to occur earlier than planned (event C) in the preventive maintenance schedule 33808. An additional maintenance event for the machine 2 33812 may be set to occur soon after the upcoming scheduled preventive maintenance event (again event C) based on, for example timing of failure of machines 3 and 4 after preventive maintenance event B. By setting a shorter interval between preventive maintenance event C and predictive maintenance event 2 (33812), a risk of a bearing-related failure may be reduced.

[1852] In example embodiments, an industrial machine predictive maintenance system may apply machine learning and the like to a range of factors to facilitate predicting and facilitating service, such as determining a schedule for service, identifying at least one qualified party for performing the service, recommending one or more sources of materials required for the service, fulfilling procurement and delivery of the materials required for the service, and rating the service of one or more parts of the industrial machine. The machine learning capability of such a system may take input, such as in the form of diagnostic-related information for the industrial machine from one of a plurality of industrial machine-related diagnostic test data, including without limitation at least one of infrared thermography of one or more parts of the industrial machine, ultrasonic testing of one or more parts of the industrial machine, motor testing of one or more parts of the industrial machine, magnetic field testing of the motor of one or more parts of the industrial machine, electron magnetic flux (EMF) testing of one or more parts of the industrial machine (e.g., pulse detection and the like), current and/or voltage testing of one or more parts of the industrial machine (e.g., from machine resident testing equipment and/or externally applied testing equipment and the like), torsional testing of one or more parts of the industrial machine (e.g., using EMF and the like), non-destructive testing of one or more parts of the industrial machine, (e.g., as may be mandatory for nuclear and power industries and the like), x-ray testing of one or more parts of the industrial machine (e.g., turbine blades and the like), video analysis for detection of vibration of one or more parts of the industrial machine, electronic field testing of one or more parts of the industrial machine, magnetic field testing of one or more parts of the industrial machine, acoustic detection of one or more parts of the industrial machine, power and/or current and/or voltage testing of one or more parts of the industrial machine, (e.g., applying algorithms comparable to those used for vibration analysis to determine when current changes are anomalies), spectrum analysis of power

consumed by a machine (e.g., a rotating machine and the like), correlation of mechanical and power faults of one or more parts of the industrial machine, sound meter for validating sound produced by or at least in proximity to one or more parts of the industrial machine, and the like. In example embodiments, machine learning may be applied to any of these sources of testing data individually to detect patterns, and the like that may be useful in detecting when a noticeable change in, for example, a detected pattern has occurred or is about to occur.

[1853] In example embodiments, combinations of diagnostic testing, such as those described herein may be used by machine learning to validate or repudiate one or more potential sources as producing anomalies that may indicate a need for service and the like. In example embodiments, combining infrared thermography with motor testing for example, such as by applying a test load onto the motor while capturing infrared images may be useful in determining combinations of conditions may indicate a potential failure, or at least a condition associated with a failure, a need for service, and the like. In example embodiments, combining, for example sounds meter capture with non-destructive testing may produce sound patterns that may be compared to baseline sounds for the specific non-destructive test condition; thereby allowing for multi-modal assessment of results (non-destructive testing results and sound test results). In example embodiments, variations in sound produced by or proximal to an industrial machine may indicate a potential failure conditions, validate a candidate failure condition, and/or diminish the likelihood of a potential failure. In example embodiments, combining multiple modes of non-destructive testing, such as acoustic and x-ray may help determine if a condition that may be detected in one of the testing modes (e.g., acoustic) correlates to a potential anomaly detectible in the other testing mode (e.g., x-ray) and the like. In example embodiments, machine learning may develop an array of test conditions, test results, and degrees of compliance with expected results for each of the diagnostic / testing scenarios described herein, and the like. Such an array may facilitate determining when anomalies represent valid potential failure conditions.

[1854] In example embodiments, each test condition, such as those described herein may be applied and results may be captured. While a given test condition is being applied, each other test condition may be applied, thereby facilitating collection of combinations of each test condition with each other test condition. Results for each combination may be captured and represented in an array, such as the array described above. Test condition combination testing may be performed when a service call, such as preventive maintenance or repair is required. In example embodiments, the industrial machine predictive maintenance system may facilitate coordinating maintenance, such as replacement of worn bearings in an industrial machine. The test condition combination array may be consulted to determine which test conditions might be applied in combination with post bearing replacement testing, such as be detecting one or more cells in the array along post bearing replacement testing axis has little or no combination data. A work order and/or procedure for post bearing replacement testing may be adapted, such as conditionally, and for specific instances, to include applying the additional testing condition indicated by the specific cell in the array. Such as approach may increase testing data, while distributing the burden of testing across time, or at least across instances of performing service on the industrial machine.

[1855] In example embodiments, machine learning may also be applied to combination condition testing, such as for detecting which combinations of testing conditions correlate best to actual failures. By learning which combinations correlate to failures, combinations that are less likely to yield a potential failure may be deprioritized so that valuable testing resources, such as service personnel and the like can be directed to combination testing with a greater likelihood of yielding actionable information.

[1856] In example embodiments, test results from a first mode of testing of a specific industrial machine, such as motor testing may be processed with machine learning algorithms and the like that may correlate certain machine testing results with one or more candidate failure modes. Test results from a second mode of testing of the specific machine, such as torsional testing may be processed with the machine learning algorithms and the like that may correlate certain torsional testing results with one or more candidate failure modes. The one or more candidate failure modes from the machine testing may be compared with those of the torsional testing. Any candidate failure modes that match for the two types of testing may be candidates for processing combined test results with machine learning. When the machine testing results and the torsional testing results are combined and processed with machine learning, candidate failure modes may be correlated thereto. If one of the candidate failure modes of the combined testing matches any candidate failure modes of the combined testing, a likelihood of the combined testing indicating a likelihood of failure may be strengthened. When such confirmation is detected through this combined testing result machine learning process, a service/repair action may be initiated to prevent failure of the specific industrial machine. In addition, testing procedures may be adapted to include combination testing so that the likely combined test result failure mode may be avoided in other industrial machines.

[1857] Referring to Figure 193, an industrial machine predictive maintenance system 33902 may execute machine learning algorithms 33904 and the like on data from a range of diagnostic testing systems, including without limitation an infrared thermography system 33906, an ultrasonic testing system 33908, a motor testing system 33910, a current and voltage testing system 33912, a torsional testing system 33914, a non-destructive testing system 33916, power, current and/or a voltage testing system 33918, a sound testing system 33920, and the like. The industrial machine predictive maintenance system 33902 may access a library of testing results 33922 that may include test results for these testing systems for prior invocations of tests on a specific industrial machine, and or on similar industrial machines. These results may be processed by the machine learning algorithms with failure mode information for the specific industrial machine and/or similar industrial machines to determine test conditions, and in particular combination of test conditions may correlate to specific failure modes. The machine learning algorithms 33904 may use artificial intelligence techniques to determine patterns, similarities, and the like among data from the library, thereby facilitating detection of combinations of testing conditions that may correlate to one or more failure modes.

[1858] In example embodiments, a method of improving correlation between diagnostic test results and machine failures may include improving correlation between results of a plurality of

diagnostic tests performed on industrial machines and failure information for failures of similar industrial machines by detecting at least one of patterns in the diagnostic test results that correlate to machine failures, similarities of diagnostic test results with machine failures. In example embodiments, a single type of machine failure correlates to failure results of a subset of the diagnostic tests.

[1859] In example embodiments, improved methods and systems for industrial machine maintenance, including methods and systems that facilitate collecting, discovering, capturing, disseminating, managing, and processing information about industrial machines, including factual information (such as about internal structures, parts and components), operational information and procedural information, including know-how and other information relevant to maintenance, service and repairs may include methods for rating a range of services and service providers associated with industrial machine predictive maintenance and the like. In example embodiments, service providers for performing maintenance and related activities may be rated. While performing a service prescribed in a service procedure, a service provider (e.g., a technician and the like) may be evaluated for the degree to which (s)he follows the procedure. The degree to which the procedure is followed may be captured implicitly by independently determining if a step has been completed in the order specified. In example embodiments, a procedure that requires removing a bearing cover panel followed by taking a photograph of the bearings may be verified by requiring the service technician to submit a photograph of the uncovered bearings before proceeding through the process. In example embodiments, the service technician may use a user interface of a computing device, such as a tablet, portable phone, industrial portable computer and the like via which the technician accesses the service procedure. The service technician may be rated along a range of criteria, including without limitation, ease of scheduling, degree of expertise/training with a specific machine and/or service activity, a result of post-service diagnostic testing (e.g., self-testing and the like), estimated versus actual costs for the service, promptness for performing the service as scheduled, cleanliness however subjective that criteria may be, adherence to procedure (e.g., as described above and the like) dependence on other resources, such as third-parties and the like.

[1860] In example embodiments, a vendor rating system 34000 is depicted in Figure 194. The vendor rating system 34000 may include a vendor rating facility 34002 that captures information about a vendor 34006 (e.g., location(s), user feedback, and the like), service data for one or more procedures 34008 that the vendor 34006 alleges to know, vendor rating weighting data 34010 that may impact how information is used to rate vendors (e.g., older data may be weighted less heavily than newer data, service on machines with very little service information may be weighted less heavily, and the like). The vendor rating system facility 34002 may further consider overall experience level of a vendor by applying an experience scale 34012 that impacts a confidence factor of a specific vendor rating based on the vendor's experience and extent of rating. Service technician input 34014 may be considered, such as structured (e.g., multiple choice responses) and/or freeform input that a service technician may provide about a service activity and the like to explain why a procedure was not followed or why a service took longer than anticipated and the

like. The vendor rating facility 34002 may further receive information from the diagnostic testing 34022, such as tests performed and results of tests associated with a service action that may be used to evaluate success of the service action performed. The diagnostic testing information 30222 may include information from diagnostics tests such as, infrared thermography, ultrasonic testing, motor testing, current/voltage testing, torsional testing, non-destructive testing, power density testing, sound testing and the like. In example embodiments, the vendor rating facility 34002 may rate vendors on a range of vendor rating criteria 34016 including, without limitation results of post service diagnostics as may be determined from the diagnostics test results data 30222 and the like. Vendor rating criteria may further include east of schedule, degree of experience with a procedure, machine, and the like, cost, promptness, cleanliness, adherence to procedures, and the like. Vendor rating results may be stored and accessed in a vendor rating results data store 34022 that may be processed with machine learning algorithms 34024 to improve correlation between, for example, a vendor rating criterion (e.g., degree of experience) and a vendor's ratings.

[1861] In example embodiments, a method of vendor rating may include determining a rating for an industrial machine service provider by gathering feedback about industrial machine services provided by the service provider and comparing the feedback to a plurality of rating criteria comprising results of diagnostics tests performed after completion of at least one industrial machine service, scheduling the service provider, cost of the service provided, promptness of the service provider, cleanliness of the service provider, adherence to a procedure for the at least one industrial machine service, a measure of experience of the service provider with at least one of the procedure and the industrial machine. In example embodiments, the method may include improving correlation of vendor rating results with rating criteria by applying machine learning to vendor rating results and incorporating an output of the machine learning when rating a vendor.

[1862] In example embodiments, improved methods and systems for industrial machine maintenance, including methods and systems that facilitate collecting, discovering, capturing, disseminating, managing, and processing information about industrial machines, including factual information (such as about internal structures, parts and components), operational information and procedural information, including know-how and other information relevant to maintenance, service and repairs may include methods for rating a range of activities and information associated with industrial machine predictive maintenance and the like. In example embodiments, procedural information for performing maintenance and related activities may be rated. While performing a service prescribed in a service procedure, a service provider (e.g., a technician and the like) may indicate a rating for each procedure, such as for each substantive service procedure action, through a user interface via which the technician accesses the service procedure. The service technician may rate each procedure along a range of criteria, including without limitation, ease of access to the information, educational value of the information, accuracy of the descriptions, accuracy of the images, accuracy of the sequence, degree of difficulty to perform the service, and the like. Service providers and the like who rely on procedural information for performing maintenance and the like on one or more machines may develop know how regarding servicing systems using such procedural information. This know how may be captured in a procedure rating system through free

form comments associated with the procedure, via suggested edits to the published procedures, and the like.

[1863] In example embodiments, a procedure to perform a maintenance task may be clear to a service technician who is familiar with the particular machine, yet it may not be sufficiently clear to service personnel with less experience. Therefore, information about the service technician completing the procedure rating task may be applied to better weight the ratings. Additionally, a service procedure may be rated on an experience scale that may facilitate identifying when a less experienced person could be used to perform a service task and when an experienced provider is preferred. Such information may be useful to an industrial machine predictive maintenance system for facilitating selection of a service entity suitable for performing a required service task and the like. In example embodiments, an industrial machine predictive maintenance system may gather information that may be descriptive of various aspects of a service / maintenance procedure, such as the experience scale rating when facilitating access to vetted service personnel. In particular, if a service procedure is rated as highly complex to follow, then service entities that have few or no experienced personnel available for performing the service may by bypassed or at least may be presented below service entities that have greater experience, greater numbers of available experienced service technicians and the like. Rating procedural information may further enhance systems for generating service procedural information by identifying characteristics of service procedures that are preferred over those that are found to be lacking and the like.

[1864] In example embodiments, such as shown in FIG 195, methods and systems for rating industrial machine service and/or repair procedures may include a procedure rating facility 34102 that may aggregate various sources of procedure rating content and produce one or more ratings for the procedure, such as ease of use, accuracy, flexibility and the like. Such a rating facility 34102 may have access to the procedure 34106, such as to process the text, images, flow charts and the like in the procedure; thereby facilitating rating various elements that contribute to the procedure. The procedure rating facility 34102 may also have access to service data 34108 for the procedure, such as a long of instance of use of the procedure, and the like. Such service data may be useful in determining a degree of confidence of a rating of the procedure. Rating for procedures that are used less often may have lower confidence than ratings for often used procedures, due at least in part to the lack of comparative data for the lower-use procedures. Rating procedures may also include accessing weighting 34110 of factors that contribute to the rating, such weighting may be explicitly stated, implicitly determined, and may vary based on factors such as age of the procedure, availability of materials required to follow the procedure, and the like. In example embodiments, rating some procedures may be impacted by experience of contributors to the rating process, such as service technicians, supervisors, procedure quality testers, and the like. Therefore, an experience scale 34112 may be applied to the rating algorithm to, for example, impact the aspects of a procedure that a contributor with given experience may be permitted to evaluate, and the like. In example embodiments, service technician and other contributor inputs 34114 to the rating process may be gathered explicitly, such as through a contributor marking a rating scale for various aspects of the procedure (e.g., the text of the procedure, the translation of a procedure, and the like).

Contributor input may be gathered implicitly, such as by tracking the time that it takes to perform the steps in the procedure, and the like. In example embodiments, if a service technician followed different steps or additional steps than those presented in the procedure, the procedure rating facility may take this input and reasons for these other steps as influence of the rating of the procedure. This feedback may help identify procedures with inaccurate machine analysis and or manufacturers guidance that may help in improving service quality. Improper machine fault diagnosis may be analyzed by artificial intelligence, such as the machine learning facility 34124 to improve analysis. Feedback from technicians and procedure rating analysis and results may be made available or pushed to the procedure developer (e.g., the industrial machine manufacturer and the like) to facilitate improving the procedure to achieve better and faster repairs. Through incentivized feedback programs and proper use thereof, such as for the rating procedures 34102, institutional knowledge may permeate every aspect of a preventive maintenance system without requiring one-on-one training like in the past.

[1865] In example embodiments, a procedure rating facility, such as the rating facility 34102 may further have access to rating criteria 34116, which may include without limitation, ease of accessing the procedure, ease of translating the procedure, educational value of the procedure, accuracy of the text, accuracy of the images/graphics, accuracy of related content (e.g., parts lists), validity of the sequence of steps, degree of difficulty overall to obtain an error free result from the procedure when using it for the first time, dependence on other steps that may or may not be directly documented, and the like. A rating facility, such as the procedure rating facility 34102 may produce procedure rating results 34122 that may be stored electronically, such as in a non-volatile computer-accessible memory and the like. In example embodiments, ratings for procedures for a specific industrial machine may be stored in one or more of the smart RFID components disposed with the machine. The procedure rating results 34122 may be improved through use of the machine learning 34124 that works cooperatively with the procedure rating facility 34102, and the like.

[1866] In example embodiments, a method for rating an industrial maintenance procedure may include determining a rating for an industrial machine service procedure by gathering feedback about the procedure from service providers who use the procedure to perform an industrial machine service and comparing the feedback to a plurality of rating criteria comprising ease of access of the procedure, ease of translation, educational value, accuracy of content, sequence accuracy, ease of following the procedure, and dependence on non-procedure actions. The method may further include improving correlation of procedure rating results with rating criteria by applying machine learning to procedure rating results and incorporating an output of the machine learning when rating a procedure.

[1867] In example embodiments, Blockchain™ techniques and applications, such as decentralized voting, cryptographic hashing, verifiability, security, open access, speed of access and update, as well as ease of adding participants (e.g., contributors, verifiers and the like) may be applied to the industrial machine predictive maintenance methods and systems described herein. Collection of data, such as operational, test, failure, and the like from industrial machines may be processed in a Blockchain™ approach that facilitates ensuring verifiability of information regarding system

status, failures, and the like. Transactions for parts orders, service orders, and the like may be processed in a Blockchain™ thereby increasing security and verifiability of transactions, including information such as costs, and the like that may be utilized by the predictive maintenance systems described herein to manage industrial machine maintenance and service activities. Other uses of block chain may include securing a distributed public ledger, such as the distributed ledger 33302 depicted in and described in association with Figure 187 herein.

[1868] In example embodiments, transactions conducted over a peer-to-peer network of industrial machines, such as IoT devices and the like may be operated as a Blockchain™ enabled distributed ledger, thereby reducing a dependency on a centralized control or repository of industrial machine and the like preventive maintenance data. In an example of Blockchain™ functionality in an industrial machine predictive maintenance system, changes to smart RFID elements on individual machines and their counterpart network-resident copy may be processed through a Blockchain™ distributed ledger system that facilitates open access to information in the RFID, such as by accessing the relevant information in the network-resident copy.

[1869] In example embodiments, FIG 196 depicts a Blockchain™ for transactions associated with a specific industrial machine 34200 that may be initiated 34202 when the industrial machine is shipped or finalized for shipment. As further transactions of the specific industrial machine are performed, such as during an installation 34204, collecting operational information from sensors deployed with the industrial machine 34206, service events of the machine 34208, parts and service orders 34210, diagnostic activity 34212, and the like each may be added to the Blockchain™ for the specific industrial machine; thereby providing a secure, verifiable, traceable data set for the industrial machine that can be leveraged by the predictive maintenance methods and systems described herein.

[1870] In example embodiments, a method of accumulating information about an industrial machine may include initiating a blockchain of industrial machine information for a specific industrial machine by generating an initiating block, and generating subsequent blocks of the specific industrial machine blockchain by combining data from at least one of shipment readiness, installation, operational sensor data, service events, parts orders, service orders, and diagnostic activity and a hash of the most recently generated block in the blockchain.

[1871] In example embodiments, predictive maintenance schedules, actions, and the like may be based on analysis of industrial machine operational data, such as data from sensors deployed with the industrial machine. Determining a maintenance triggering threshold for operational data, including sensed data, may include identifying a type of effect the data represents and then determining data values that represent acceptable operation, questionable operation, unacceptable operation, and other types of operation. In example embodiments, vibration sensors deployed to detect and monitor vibration activity of industrial machine components, structural elements, and the like may facilitate determining how vibration of machine parts contributes to predictive maintenance actions. Determining a severity of vibration data from the sensors relative to timing and the like of predictive maintenance actions may require more than conventional vibration

analysis. In example embodiments, vibration measures may be translated into severity units that may be used when predicting maintenance requirements and the like.

[1872] In example embodiments, while vibration may be useful for determining negative effects on industrial machines, vibration analysis is generally complex and varies greatly based on frequency of vibration, vibration source, material being vibrated, machine operating cycles per minute, and the like. A measure of vibration, such as vibration velocity may be useful for determining when vibration is a problem for a mid-range vibration frequency, but alone it fails to usefully provide insight at low and high frequencies. Therefore, vibration analysis that is frequency independent, such as vibration analysis measures that are normalized, may result in useful predictive maintenance information.

[1873] In example embodiments, normalizing vibration analysis results into severity units as described herein may facilitate vibration frequency independence. Overall vibration spectra, RMS levels, and the like may be expressed in units of displacement, velocity, acceleration and the like. In an example, bearing cap vibration readings may be expressed as vibration velocity at least because it directly relates to mechanical severity of the vibration. As noted above, while vibration velocity may be sufficient for mid-range frequency components, low and high frequency components exhibit significant exceptions to the relevance of vibration velocity for predictive maintenance algorithms. It will be appreciated in light of the application that vibration velocity man be characterized through amplitude-versus-frequency charting and the like that, in effect, linearly lower the velocity severity requirements (e.g., vibration amplitude and the like) for low and high frequencies, such as when compared to mid-range frequency velocity severity requirements.

[1874] In example embodiments, the methods and systems described herein extend and enhance methodologies of frequency charting to facilitate normalizing vibration spectra so that it can be expressed as vibration severity units that are consistent across wide vibration frequency spectra, such as from near-zero frequency to well over 18,000 cycles per minute (cpm). Components of the vibration spectra that occur at frequencies below a low-end linearity frequency (e.g., a low-end knee frequency value) will be processed with an algorithm that normalizes to a value of displacement (e.g., a preset value of millimeters of displacement) because displacement (e.g., amplitude) has been shown to be a more significant indicator of severity than velocity at lower frequencies. Components of vibration spectra that occur at frequencies above a high-end linearity frequency (e.g., a high-end knee frequency value) will be processed with an algorithm that normalizes to a value of units of gravity (e.g., a preset value of g's or g force). The net result is that each range of the frequency spectra (below the low-end knee threshold, mid-range, and above the high-end knee threshold) can be mapped uniformly to severity units. In many examples, the frequency spectra may be broken into three ranges (below low-end knee threshold, mid-range, and above high-end knee threshold), fewer or more ranges of frequency spectra may be determined and applied without exceeding the scope of the vibration data normalization techniques for generating predictive maintenance vibration severity units.

[1875] In example embodiments, methods and systems include normalizing vibration amplitude units into units that are independent of frequency. These units can be referred to as severity units or action units. In many examples, vibration spectra, overall levels or root-mean-square levels are expressed in units of displacement, velocity or acceleration. For bearing cap readings, for example, vibration velocity is most commonly used as it may be directly related to mechanical severity. Although sufficient for mid-frequency components, there can be, however, significant exceptions for low frequency and high frequency domains. It will be appreciated in light of the disclosure that many amplitude-versus-frequency severity charts have been constructed to linearly lower the velocity severity requirement for both the lower and the higher frequency components depicted in the chart.

[1876] In example embodiments, the methods and systems include development and construction of a severity graph to normalize vibration spectra as severity units. By way of this example, lower frequency components below a predetermined knee level of about 1,200 cycles per minute, as depicted in Figure 176, will be gained by a predetermined factor (as a function of the slope) such that its amplitude in severity units may be normalized with respect to severity. Similarly, for higher frequency components above a knee level of about 18,000 cycles per minute, spectral peaks are also gained by a different predetermined factor to achieve severity flatness. In example embodiments, spectra displayed in severity units may be displayed with horizontal lines to demarcate severity. In many aspects of the embodiments, other spectral components related to one or more bearing defect frequencies and/or one or more bearing resonance frequencies may have their corresponding amplitudes adjusted for severity. By way of this example, other spectral components related to one or more bearing defect frequencies may have their corresponding amplitudes increased to adjust for severity, other spectral components related to one or more bearing resonance frequencies may have their corresponding amplitudes decreased to adjust for severity. In addition, other digital processing techniques, which output spectra such as enveloping, may be employed to supplement or superimpose spectral peaks within the severity spectrum. In example embodiments, the final resulting severity spectrum may then be displayed local, remotely and/or accessed through a cloud network facility for presentation and analytical purposes. In example embodiments, the final resulting severity spectrum may be fed to an expert system for analysis and evaluation of the severity. In many aspects of the embodiments, an overall level may be calculated or derived from this "normalized" spectrum to produce an overall level or a root-mean-square level in units of severity rather than the more typically collection of disparate units currently utilized by vibration monitoring systems.

[1877] In example embodiments, Figure 197 depicts a diagram showing a severity unit conversion function for normalizing vibration sensor data for casing vibration on industrial machinery. The severity unit conversion function 30602 includes vibration displacement rate (inches per second) along a vertical axis 30604 and vibration frequency cpm (cycles per minute) along a horizontal axis 30606. A low-end frequency demarcation 30608 is set at 1200 cpm, defining the upper end of the low-end vibration frequency region 30610 as well as the lower end of the mid-frequency region 30612. A high-end frequency demarcation 30614 is set at 18000 cpm, defining a lower end of the

high-end vibration frequency region 30616 as well as the high-end of the mid-frequency region 30612.

[1878] Severity for the embodiment of Figure 176 is calculated as follows:

[1879] S=MxA (30601)

[1880] In the equation 30601, S is the severity value being calculated, A is a mid-range severity limit, and M is a severity normalizing value that is calculated for each of the three vibration spectra ranges as follows: for the low-end range 30610: M= vibration frequency/low-end demarcation value; for the mid-range 30612: M=1; and for the high-end range 30616: M=high-end demarcation value/vibration frequency.

[1881] In the example of the embodiments of Figure 197, the low-end range M=frequency/1200 and for the high-end range M=18000/frequency. This results in an acceptable severity value of approximately 2.5 mils for the low-end range and 2.5 g's for the high-end range.

[1882] In example embodiments, the severity normalization function exemplified in Figure 197 can facilitate developing severity units for each frequency range that may be used by the predictive maintenance methods and systems described herein.

[1883] In example embodiments, five severity units are identified and may be applied to each frequency range. Severity units may be named: acceptable, watch, resurvey, action soon, immediate, and the like. In example embodiments, vibration data that results in an acceptable severity unit has little, if any, impact on predictive maintenance analysis and action recommendations. Vibration sensor data studies that result in acceptable severity unit analysis may be gathered and further analyzed for variations among industrial machines, such as similar industrial machines, similar portions of industrial machines, different generations of industrial machine or portion thereof and the like.

[1884] In example embodiments, additional severity categories may be added as depicted in Figure 198. With continuing reference to Figure 198, the exemplary severity chart may define severity levels with associated actions for those levels. By way of this example, the severity chart may be associated with spectral peaks taken with a bearing cap mounted accelerometer. The range at which the one or more detected signals are deemed acceptable and, therefore, the least severe across the three ranges of the detected signal are less than about 2.5 thousandths of an inch peak-to-peak (about 63.5 micrometers peak-to-peak) when measuring displacement for a regime that is less than about 1,200 cycles per minute or less than about 20 Hz. For the regime that is about 1,200 cycles per minute to about 18,000 cycles per minute or about 20 Hz to about 300 Hz, the severity chart may assess signals in terms of velocity and the acceptable and, therefore, least severe level is less than about 0.15 inches per second at peak (about 3.81 millimeters per second at peak). For the regime that is greater than about 18,500 cycles per minute or greater than about 300 Hz, the severity chart may assess signals in terms of acceleration and the acceptable and, therefore, least severe level is less than about 2.5 g level at peak.

[1885] The range at which the one or more detected signals are deemed worthy of watching and, therefore, one level higher than the least severe across the three ranges of the detected signal are between 2.5 thousandths of an inch peak-to-peak (about 63.5 micrometers peak-to-peak) and 5

thousandths of an inch peak-to-peak (about 127 micrometers peak-to-peak) when measuring displacement for a regime that is less than about 1,200 cycles per minute or less than about 20 Hz. For the regime that is about 1,200 cycles per minute to about 18,000 cycles per minute or about 20 Hz to about 300 Hz, the severity chart may assess signals in terms of velocity and the worth to watch and, therefore, one level higher than the least severe level is between about 0.15 inches per second at peak (about 33.8 millimeters per second at peak) and about 0.3 inches per second at peak (about 67.6 millimeters per second at peak). For the regime that is greater than about 18,500 cycles per minute or greater than about 300 Hz, the severity chart may assess signals in terms of acceleration and the worthy to watch and, therefore, one level up from the least severe level is between about a 2.5 g level at peak and about a 5 g level at peak.

[1886] The range at which the one or more detected signals are determined to be sufficient to suggest or require a re-survey of the machine or route from which the one or more signals were obtained and, therefore, one level higher in severity than the watch level and two levels of severity higher than the least severe across the three ranges of the detected signal are between 2.5 thousandths of an inch peak-to-peak (about 63.5 micrometers peak-to-peak) and 5 thousandths of an inch peak-to-peak (about 127 micrometers peak-to-peak) when measuring displacement for a regime that is less than about 1,200 cycles per minute or less than about 20 Hz. For the regime that is about 1,200 cycles per minute to about 18,000 cycles per minute or about 20 Hz to about 300 Hz, the severity chart may assess signals in terms of velocity and define a range in which it may be sufficient to suggest or require a re-survey of the machine or route from which the one or more signals were obtained between about 0.3 inches per second at peak (about 7.62 millimeters per second at peak) and about 0.6 inches per second at peak (about 15.24 millimeters per second at peak). For the regime that is greater than about 18,500 cycles per minute or greater than about 300 Hz, the severity chart may assess signals in terms of acceleration and be sufficient to suggest or require a re-survey of the machine or route from which the one or more signals were obtained between about a 5 g level at peak and about a 10 g level at peak.

[1887] By way of this example, the range at which the one or more detected signals are determined to be sufficient to flag for action soon and, therefore, one level below a severity level to flag for action. In other examples, there can be a flag for action now and a flag action including a flag for shutdown when the severity of one or more detected signals warrant such a flag. When measuring displacement for a regime that is less than about 1,200 cycles per minute or less than about 20 Hz, the sufficient to flag for action soon range may be between about 10 thousandths of an inch peak-to-peak (about 254 micrometers peak-to-peak) and about 16.6 thousandths of an inch peak-to-peak (about 421.64 micrometers peak-to-peak). For the regime that is about 1,200 cycles per minute to about 18,000 cycles per minute or about 20 Hz to about 300 Hz, the severity chart may assess signals in terms of velocity and define a range in which it may be sufficient to suggest or require a re-survey of the machine or route from which the one or more signals were obtained between about 0.6 inch per second at peak (about 15.24 millimeters per second at peak) and about 1 inch per second at peak (about 25.4 millimeters per second at peak). For the regime that is greater than about 18,500 cycles per minute or greater than about 300 Hz, the severity chart may assess signals

in terms of acceleration and be sufficient to suggest or require a re-survey of the machine or route from which the one or more signals were obtained between about a 10 g level at peak and about a 16.6 g level at peak.

[1888] By way of this example, the range at which the one or more detected signals are determined to be sufficient to flag for immediate action and, therefore, at the highest severity level. In other examples, there can be a flag for immediate action and a flag action including a flag for shutdown when the severity of one or more detected signals warrant such a flag. When measuring displacement for a regime that is less than about 1,200 cycles per minute or less than about 20 Hz, the sufficient to flag for immediate action soon range may be above about 16.6 thousandths of an inch peak-to-peak (about 421.64 micrometers peak-to-peak). For the regime that is about 1,200 cycles per minute to about 18,000 cycles per minute or about 20 Hz to about 300 Hz, the severity chart may assess signals in terms of velocity and define a range in which it may be sufficient to flag for immediate action above about 1 inch per second at peak (about 25.4 millimeters per second at peak). For the regime that is greater than about 18,500 cycles per minute or greater than about 300 Hz, the severity chart may assess signals in terms of acceleration and be sufficient to flag for immediate action soon above about a 16.6 g level at peak.

[1889] It will be appreciated in light of the disclosure that the severity chart in Figure 197 depicts 0.15 inch per second velocity at 1,250 cycles per second in the Acceptable category. The conversion between displacement, velocity and acceleration depicted in Figure 197 shows that 2.5 thousandths of an inch displacement peak-to-peak is equivalent to 0.15 inches per second velocity at 1,250 cycles per second in the normalization to determine severity units. Figure 197 also shows that 0.2 inches per second velocity at peak at 61,450 cycles per minute is equivalent 2.5 g level of acceleration. The Watch category spans 6 dB. The Resurvey category spans 6 dB and the Action Soon category spans about 4.5 dB.

[1890] It will be appreciated in light of the disclosure that many examples of severity charts may be based on highly specific equipment types. In many examples, some of these classifications may be simplified because many categories of machines that run at sufficiently low or relatively slower speeds may not need separate severity categories. In these examples, severity units based on velocity may be sufficient to provide one or diagnoses. In many examples, communication between different subsystems such as a raw data server that may serve up vibration waveform, spectrum and overall levels and an expert system engine that must translate this raw data into meaningful severity units may be significantly simplified by the use of normalizations to produce the severity units.

[1891] In example embodiments, the severity units may be applied to non-vibration data where signal processing techniques may be applied to any raw set of data that has specialized significance, but which must be normalized to be successfully compared or analyzed. In example embodiments, actuarial data regarding the viability of a specific pharmaceutical treatment that may be gender specific may be normalized to the general population. It will be appreciated in light of the disclosure that one or more established techniques or guidelines normalizing the gender-specific

data to a gender-less universe becomes useful for subsystem communication to AI, statistical, tutorial or other relevant systems.

[1892] In example embodiments, vibration data that results in a watch severity unit may impact aspects of predictive maintenance recommendations, such as a frequency of occurrences of vibration data collection and analysis. Watch severity unit determination may result in conducting at least vibration data collection and analysis more frequently. It may also result in checking other conditions of the components being vibrated, such as by performing calibration, diagnostic testing, visual inspection and the like.

[1893] In example embodiments, vibration data that results in a resurvey severity unit may trigger performing vibration data collection and analysis as soon as possible. Resurvey severity unit determination may result in a signal (e.g., a set of commands and the like) being transmitted to relevant portions of the affected industrial machine to configure the data collection and routing functionality and elements to repeat the vibration data collection and analysis again. It may also result in configuring the industrial machine data collection control systems to initiate data collection from other sensors for the involved industrial machine elements. Likewise, it could raise the priority of collecting comparable vibration sensor data from other similar industrial machines so that it can be available for comparative analysis of the resurveyed vibration study and the like.

[1894] In example embodiments, vibration data that results in an action soon severity unit may trigger scheduling a service action of the affected parts well ahead of a next scheduled maintenance for a portion of the industrial machine with the affected parts. It may also result in escalating actions (e.g., preventive, survey, analysis, and the like) for related elements. In an example, if vibration data for a motor indicates taking action soon, vibration data collection, preventive maintenance actions, calibration actions and the like may be activated for a drive shaft of the motor, a gearbox being driven by the driveshaft, and the like.

[1895] In example embodiments, vibration data that results in an immediate severity unit may be treated as constructive approval to perform all necessary part replacement as soon as possible, thereby triggering ordering of replacement parts, materials, and the like to perform one or more service actions on the industrial machine. Such a result may also trigger certain automatic actions such as stopping use of the industrial machine, reducing the duty cycle of the industrial machine, reducing an operating cycle rate of the industrial machine, and the like until service is performed, and the like.

[1896] An embodiment of severity units applied to vibration across a wide vibration frequency range is representatively depicted in Figure 198. In the representative embodiment of Figure 198, each of five severity units are mapped to the three vibration spectra regions represented in Figure 197, specifically for vibration frequencies below 1200 cpm, between 1200 cpm and 18000 cpm, and above 18000 cpm.

[1897] In example embodiments, within each spectral region severity units are defined. For the spectral region below the low-end threshold (e.g., 1200 cpm), vibration displacement below 2.5 mils peak-to-peak meets the acceptable severity unit criteria; between 2.5 and 50 indicates a watch severity unit; between 5.0 and 10.0 indicates a resurvey severity unity; between 10.0 and 16.6 mils

displacement indicates an action soon severity unit, and displacement greater than 16.6 mils triggers an immediate action severity unit. For vibration frequency spectra between 1200 cpm and 18000 cpm, normal severity is characterized by displacement below .15 inches per second peak (ipsp); watch is between 0.15 and 0.3 ipsp; resurvey is between .3 and .6 ipsp; action soon severity occurs between 0.6 and 1.0 ipsp; and immediate action severity occurs for vibration displacement rates greater than 1.0 ipsp. For vibration frequency spectra greater than 18000 cpm, acceptable severity is indicated by vibration analysis indicating less than 2.5 gs peak; watch is indicated by 2.5 gs to 5.0 gs; resurvey for 5.0 gs to 10.0 gs; action soon for 10.0 gs to 16.6 gs; and immediate action severity unit is indicated for vibration that results in forces greater than 16.6 gs.

[1898] Applications of the severity unit methods and systems described herein include uses across a range of machines operating at various speeds. Unlike existing vibration analytical tools, the algorithm-based approach described herein can readily handle slower speed machines by effectively removing some unnecessary computational complexity associated with an impact of machine speed, and the like. In environments where different machines perform different actions, such as raw data analysis and severity detection, communication bandwidth must be increased to support providing enough information to ensure robust severity determination. Use of the severity unit methods and systems described herein significantly simplify data communication needs in such embodiments; thereby reducing communication bandwidth demand in corresponding environments and the like.

[1899] While this discussion of severity units is directed at vibration data analysis and the like, the methods and systems for severity unit determination and detection may be applied to data sources other than vibration that can benefit from normalization for successful comparison. In example embodiments, actuarial data regarding the viability of a specific pharmaceutical treatment for one or both genders may be normalized using the methods and systems described herein to be applied to the general population. Algorithms may be generated that accommodate existing guidelines for severity, yet extend them using the methods and systems described herein to produce gender-less (gender normalized) severity measures.

[1900] In example embodiments, a method of predicting a service event from vibration data may include a set of operational steps including capturing vibration data from at least one vibration sensor disposed to capture vibration of a portion of an industrial machine. The captured vibration data may be processed to determine at least one of a frequency, amplitude, and gravitational force of the captured vibration. Next, a segment of a multi-segment vibration frequency spectra that bounds the captured vibration may be determined, based on for example the determined frequency. Thus, calculating a vibration severity unit for the captured vibration may be based on the determined segment and at least one of the peak amplitude and the gravitational force derived from the vibration data. Additionally, the method may include generating a signal in a predictive maintenance circuit for executing a maintenance action on the portion of the industrial machine based on the severity unit.

[1901] In example embodiments, the segment is determined based on comparing the determined frequency to an upper limit and a lower limit of a mid-segment of the multi-segment vibration

frequency spectra. A first segment of the multi-segment vibration frequency spectra may include determined frequency values below a lower limit of a mid-segment of the multi-segment vibration frequency spectra. The lower limit of the mid-segment of the multi-segment vibration frequency spectra may be 1,200 kHz and the upper limit may be 18,000 kHz. In example embodiments, a second segment of the multi-segment vibration frequency spectra may include determined frequency values above an upper limit of a mid-segment of the multi-segment vibration frequency spectra.

[1902] In example embodiments, calculating a vibration severity unit may include producing a severity value by multiplying one of a plurality of severity normalizing parameters by a mid-range severity limit and mapping the vibration severity value to one of a plurality of severity unit ranges of the determined segment. A first severity normalizing value of the plurality of normalizing values is calculated by dividing the determined frequency by a low-end frequency value of the mid-segment of the multi-segment vibration frequency spectra. A specific one of the plurality of severity normalizing parameters includes the first severity normalizing value when the determined frequency value is less than the low-end frequency value.

[1903] In example embodiments, a second severity normalizing value of the plurality of normalizing values is calculated by dividing a high-end frequency value of the mid-segment of the multi-segment vibration frequency spectra by the determined frequency. A specific one of the plurality of severity normalizing parameters includes the second severity normalizing value when the determined frequency values is greater than the high-end frequency value.

[1904] Regarding segments of the multi-segment vibration frequency spectra, a first segment of the multi-segment vibration frequency spectra is divided into a plurality of severity units based on the determined amplitude of vibration. A second segment of the multi-segment vibration frequency spectra is divided into a plurality of severity units based on the determined gravitational force.

[1905] In example embodiments, the vibration severity unit is determined based on a peak displacement of the determined amplitude of vibration for determined vibration frequencies within the first segment of the multi-segment vibration frequency spectra. In an example, the vibration severity unit is determined based on the determined vibration-induced gravitational force for determined vibration frequencies within the second segment of the multi-segment vibration frequency spectra.

[1906] In example embodiments, the portion of the industrial machine may be a moving part, a structural member supporting a moving part, a motor, a drive shaft, and the like.

[1907] In example embodiments, a system for predicting a service event from vibration data may include an industrial machine that includes at least one vibration sensor disposed to capture vibration of a portion of the industrial machine. The system may further include a vibration analysis circuit in communication with the at least one vibration sensor and that generates at least one of a frequency, peak amplitude, and gravitational force of the captured vibration. The system may yet further include a multi-segment vibration frequency spectra structure that facilitates mapping the captured vibration to one vibration frequency segment of the multiple segments of vibration frequency. Also, the system may include a severity unit algorithm that receives the determined

frequency of the vibration and the corresponding mapped segment and produces a severity value which is then mapped to one of a plurality of severity units defined for the corresponding mapped segment. In example embodiments, the system may also include a signal generating circuit that receives the one of the plurality of severity units, and based thereon, signals a predictive maintenance server to execute a corresponding maintenance action on the portion of the industrial machine.

[1908] In example embodiments, the system may calculate the vibration severity level via vibration severity calculation software. The vibration severity calculation software may be configured to digitally substantially perform the functions of one or more of the vibration analysis circuit, the multi-segment vibration frequency spectra structure, and the severity unit algorithm and may be configured to be run by any general-purpose processor or otherwise suitable machine. The vibration severity calculation software may be configured to receive an input of a signal from the vibration sensor. The signal may be a digital signal or an analog signal and may include a vibration waveform, i.e., a captured vibration.

[1909] In example embodiments, the vibration severity calculation software may digitally implement one or more of high-pass filtering, low-pass filtering, integration, and differentiation of the signal received from the vibration sensor to calculate the vibration severity level. The vibration severity calculation software may generate at least one of a frequency, peak amplitude, and gravitational force of the captured vibration from the vibration sensor. The vibration severity calculation software may map the captured vibration to one vibration frequency segment of the multiple segments of vibration frequency. The vibration severity calculation software may produce the severity value based on the determined frequency of the vibration and map the severity value to one of a plurality of severity units defined for the corresponding mapped segment.

[1910] In example embodiments, the severity unit may be outputted by the vibration severity calculation software to a user or an analyst, and/or to one or more of the expert systems so that action may be taken based thereon. In example embodiments, the vibration severity calculation software may receive the one of the plurality of severity units and signal a predictive maintenance server to execute a corresponding maintenance action on the portion of the industrial machine from which the captured vibration was captured, the corresponding maintenance action being based on the one of the plurality of severity units. The vibration severity calculation software may be implemented to calculate the vibration severity level in place of or in addition to one or more of the vibration analysis circuit, the multi-segment vibration frequency spectra structure, and the severity unit algorithm.

[1911] In example embodiments, vibration-related data collected from sensors disposed with an industrial machine may include displacement, velocity, acceleration, and the like. Additionally, data such as velocity, acceleration and the like may be calculated from raw collected data, such as displacement gathered over known units of time and the like. Velocity may be based on a count of detectable vibration events in a specific period. Velocity may be independent of a size or length of a displacement occurrence. In example embodiments, acceleration may be calculated as a rate of change of velocity measures. In example embodiments, acceleration may be generated from one

or more acceleration sensors that may detect a time of a start of displacement and relative time of an end of displacement in a specific direction and based thereon may identify an acceleration of the part during a vibration occurrence. Vibration data may be helpful in determining if a part may be subject to excessive vibration. Analyzing such vibration data to make the determination involves factoring in aspects of vibration, such as frequency and the like. As described herein, conventional approaches to vibration analysis for determining a degree to which detected vibration may be unacceptable, requires evaluating vibration in different portions of the vibration spectra differently. A novel approach to normalize evaluation of an impact of vibration across an extended range of vibration spectra, such as a threshold of vibration beyond which the vibration is likely to cause a problem, such as a breakdown of the vibrating component may benefit predictive maintenance systems, such as expert systems and the like that may attempt to provide actionable information to machine owner and the like.

[1912] In example embodiments, Severity Units may facilitate normalizing vibration analysis for the purposes of determining if detected vibration is unacceptable by eliminating, or at least obfuscating the need for calculating multiple vibration measures across a range of vibration spectra. By normalizing different units of vibration measure over spectral ranges, Severity Units, also referred to herein as Action Units, may facilitate application of Severity Units for a wide range of vibration analysis applications, including without limitation, industrial machine vibration analysis, moving part vibration analysis, complex mechanical system vibration and the like.

[1913] In example embodiments, the system may normalize one or more severity units using included (or accessed) severity normalization methodologies. In example embodiments, the severity normalization methodologies may execute an envelope analysis method. In example embodiments, the severity normalization methodologies may scan a stream of vibration severity units with a band-pass filter, e.g., a band-pass filter having a width of 500 Hz, over a plurality of bands having little to no overlap, e.g., 1 kHz to 40 kHz. The severity normalization methodologies may include processing each of the scanned bands, e.g., via harmonic filtering, to analyze running speeds and electrical signals thereof to determine an envelope. With this, overall AC and DC values of the envelope can be computed and optimum regions for location of a band-pass filter based on the AC and DC values can be determined. In these examples, AC values may be used by the severity normalization methodologies to detect modulation of bearing defect frequencies. In further examples, DC values may be used to determine issues such as insufficient lubrication. By way of these examples, the determined band-pass filter location may be referred to as an envelope spectrum. In example embodiments, the severity normalization methodologies may superimpose envelope spectrums from different severity units at differing frequencies. In these examples, the severity normalization methodologies may be configured to be run by any general-purpose processor or otherwise suitable machine.

[1914] In example embodiments, the severity normalization methodologies may include the application of waveform analysis processes, such as overall, true peak, peak-to-peak, crest-factor, K-factor, product of crest-factor and amplitude. In example embodiments, the severity normalization methodologies may further include the application of statistical stability

measurement techniques to the vibration waveforms within the envelope spectrum. In these examples, the waveforms may be labeled according to results of the waveform analysis processes. In example embodiments, the severity normalization methodologies may implement phase stability spectrum analysis by marking trends in phase variation of vibration waveforms over time in a stream of severity units. In example embodiments, the severity normalization methodologies may also implement phase stability spectrum analysis by marking trends in phase variation over time of the vibration waveforms directly. In doing so, the severity normalization methodologies may include the qualification of stability of the phase variation. In example embodiments, the severity normalization methodologies may implement amplitude stability spectrum analysis (in contrast to phase stability spectrum analysis) by marking trends in amplitude variation of vibration waveforms over time in a stream of severity units and/or a vibration waveform directly. In example embodiments, the amplitude stability spectrum analysis may include the qualifying of the stability of the phase variation. In example embodiments, the severity normalization methodologies may include production of histograms of phase, amplitude, and other characteristics of vibration waveforms for analysis by users, analysts, and/or expert systems.

[1915] In example embodiments, Figure 199 depicts a vibration severity graph that charts vibration frequency along the horizontal axis. The graph includes two vertical axes – one that represents traditional vibration measures that are frequency dependent; the other represents Severity Units that are independent of frequency. The traditional vibration measures a line 30802 shows three segments, indicating safe vibration limits for three ranges of frequency. A severity units line 30804 shows a single horizontal line indicating a safe vibration-severity limit for all ranges of frequency. For traditional vibration analysis derivatives of vibration are adjusted for frequency. Such derivatives below the line 30802 may represent acceptable levels of vibration. Similarly, vibration derivatives above the 30802 may represent unacceptable levels of vibration. However, the function required to determine whether a sample of vibration results in a derivative above or below the line 30802 is different for different vibration frequencies. The knee values 30806 and 30808 may typically, as described herein align with vibration frequencies of 1,200 CPM and 18,000 CPM; however, material type, vibration object type and other factors may further impact the function to perform. In contrast, the methods and systems described herein for generating and using Severity Units and/or Action Units may be adapted to generate a normalized limit for vibration severity as represented by the line 30804. Severity/Action unit-based calculated measures of vibration below the line 30804 may indicate safe vibration limits; whereas severity/action unit-based measures above the line 30804 may represent unacceptable levels of vibration. An expert system, such as a system for predicting maintenance events for industrial machines may apply severity/action unit values for industrial machines in a simple comparison function that compares a severity/action unit value to the severity/action unit threshold value. When the unit value is below the threshold value, an impact on a prediction of a need for maintenance may be small or negligible. When the unit value is above the threshold value, an impact on a prediction of a need for maintenance may be substantive and may directly trigger predicting a maintenance event. Alternatively, the result of the comparison of a unit value with a threshold value may be used to adjust a weighting of other factors

being processed to predict a maintenance event. Through severity/action unit weighting of other factors, predicting maintenance needs for industrial machines may combine below threshold or marginal results for vibration and other factors into a prediction of industrial machine maintenance.

[1916] In example embodiments, severity units may be calculated using other signal processing techniques. These other signal processing techniques may produce an Action Unit normalized representation of the sensed vibration data. In example embodiments, other frequency thresholds may be used with various techniques and may be dependent on various factors of the machine part(s) being vibrated, such as without limitation severity peak vibration levels, gas pulse frequency peak levels, machinery component type, bearing fault frequencies and the like. In example embodiments, normalized severity/action units may be weighted based on a component type for applications, such as hammer mills, crushers, large horsepower prime movers, soft-foundation (e.g., spring isolated) and the like. While the example of Figure 178 and others in this specification use a low threshold of 1200 Hz and a high threshold of 18,000 Hz, other values can be used, such as a low threshold of 500 Hz and a high threshold of 5,000 Hz and the like. The relationship between a low threshold and a high threshold for a given application may be based on a material, operating frequency, severity sensitivity, and the like.

[1917] Vibration events that may be detected through envelop processing and the like, such as for roller bearing defects that cause machine cycle dependent vibration events (e.g., a jolt as the roller bearing impacts the defect). Once vibration events detected through envelop processing are captured, they can be processed to result in a peak value that can be mapped to a severity unit frequency spectra. In this way, envelope-detected vibration events that may be filtered out through RMS or similar time-averaging calculations, can be mapped onto a Severity/Action Unit frequency chart.

[1918] In example embodiments, severity for various components in an industrial machine or portion thereof (e.g., a gear box and the like) may be combined into an overall severity for the machine/portion. One approach is to generate an aggregated severity value by summing all the severity unit calculations for one or more components in the machine/portion. Another approach is to calculate an overall average severity for a machine/portion, such as by determining an average of the generated severity values. Other approaches for calculating an overall severity for a machine/portion may include weighting a portion of the individual component's severity value, and the like.

[1919] In example embodiments, calculations of severity units for industrial machine components, such as moving parts in an industrial machine (e.g., gears, shafts, motors, too heads, and the like) may be mapped onto a severity graph as depicted in Figure 198 and described herein, such as by identifying in the map a correspondence between a spectral peak level and a measure of severity level. A mapped severity level may be determined based on the identification. Graphical elements may be assigned to each severity level so that a severity of an industrial machine component may be presented pictorially as, for example, an overlay of an image, drawing, or other representation that shows individual components in an industrial machine. Figure 200 depicts a block diagram representing components 30902 of an industrial machine 30900 with severity unit levels indicated

by a graphical overlay elements 30904. In example embodiments, the overlay image 30904' may be presented in a graphical user interface that may facilitate data discovery by a user who interacts with the overlay by, for example touching or otherwise selecting one of the graphical overlay elements 30904. Such a scenario is depicted in Figure 200. Component severity and related information in pop-up window 30908 is visualized in response to a user selecting the graphical overlay element 30904. In example embodiments, the graphical overlay elements 30904 may represent composite severity levels for a group of components, such as a gear box, motor assembly and the like. When a composite graphical overlay element is selected, a second image, such as a detail of a gear box and the like may be visualized in the graphical user interface so that the user can dive into further details for the components in the assembly, and the like.

[1920] In example embodiments, severity units may be presented in context of a Master Action Unit Nomogram (MAUN). In example embodiments, vibration data may be collected for at least three dimensions; therefore, a 3-D MAUN that presents vibration data in action or severity units in a 3-D presentation may be produced.

[1921] In example embodiments, raw vibration data may be provided to a predictive maintenance system, such as a system that applies techniques such as machine learning and the like to determine threshold for acceptable vibration across a range of spectra. However, learning from this raw information may require information about the environment and vibration analysis engineering that results in a highly complicated maintenance prediction operation. Severity Units, such as those described herein, including MAUN and the like, may be provided to the predictive maintenance system to simplify learning by more efficiently matching raw vibration data with normalized measures of vibration severity (e.g., Severity Units and the like). Use of Severity Units and the like may further reduce filtering and evaluation complexity for predictive maintenance systems since at least some portion of these operations may be incorporated into the generation of Severity Unit measures from the raw vibration data.

[1922] In example embodiments, learning from such systems may be applied to Severity Unit calculation functions, such as may be performed locally by a data collection agent, local network processor, and the like as feedback. This feedback may be applied to threshold refinement algorithms that adjust, for example, severity level (e.g., threshold) determination from raw vibration data, so that vibration thresholds can be tuned for local conditions, and the like. Such feedback may further be useful in processes that attempt to determine which of a plurality of data processing techniques / algorithms (e.g., to produce Severity and/or Action Units and the like) may produce more accurate MAUN measures. Doing so may reduce processing complexity and reduce data storage demand, which may be desirable for reducing overall cost and sophistication of data collection devices and the like that may produce Severity Unit data.

[1923] In example embodiments, predictive maintenance methods and systems may be applied to industrial machines, such as rotating equipment machines. Exemplary rotating equipment machines for which methods and systems of predictive maintenance described herein can be used may include, without limitation drills, boring heads, polishers, motors, turbines, gear boxes, transmissions, rotary-vibratory adapters, drive shafts, computer numerical controlled (CNC)

routers, lathes, mills, grinders, centrifuges, combustion engines, compressors, reciprocating engines, pumps, fans, blowers, generators, and the like. Manufacturers of exemplary rotating equipment and related parties, such as testing services, component manufacturers, sub-contractors, and the like may have access to technical data about such equipment on a machine-by-machine basis. Additionally, information that may be available about machines, sub-assemblies, individual components, accessories, rotating integrated parts, and the like may include design parameters, test specifications, operating specifications, revisions to the products, and the like. This and related information may apply to one or more deployed machines, such as to a specific serial number, a product line of industrial machines, a given production version, a production run, and the like. Machine information available may cover aspects of the equipment that relate to one or more rotating components, such as a count of gear teeth of one or more gears (e.g., a gear box such as a helical gearbox, worm reduction gearbox, planetary gearbox and the like, a power transfer gear set, and the like), a count of motor rotor bars (e.g., rotor bars in a squirrel-cage rotor and winding, such as a synchronous motor, and the like), RPM rate for rotating components and the like. Additionally, information may be available and utilized for predictive maintenance event planning and execution of industrial machines, such as roller bearing-based systems including, without limitation (count of roller balls, count of balls, count of balls/roller, ball-to-roller contact angle(s), race dimension (e.g., inner and outer race dimensions), count of vanes, count of flutes, mode shape (e.g., relative displacement and the like) data.

[1924] Providing access to rotating equipment information, such as that exemplarily described herein, for predictive maintenance processing, such as with a predictive maintenance analysis circuit, may be automated through a range of means including, without limitation; (i) storing data that contains information about a portion of a rotating equipment machine in a non-volatile storage element integrated with or into the machine, or portion thereof, prior to deployment in the field; (ii) updating a non-volatile storage element integrated with or into the machine with the relevant rotating component information after or as part of deployment, such as during a deployment validation operation and the like; (iii) storing data representative of the rotating equipment specifications, measurements, production testing, and the like in a network accessible data storage facility (e.g., a cloud-based data storage facility indexed by at least one of part, sub-system, machine or the like identifier, such as a serial number or set thereof that associates a part (e.g., a roller bearing assembly) with a machine/deployment; (iv) a combination of (i) or (ii) and (iii), with at least a subset of information stored in the non-volatile data storage facility deployed with the machine (e.g., a serial number of the machine, serial number(s) of rotating equipment components, and the like) that can be used to identify the relevant information for a deployed machine from the network accessible data storage facility. To address commercial confidentiality concerns, some and/or all network-accessible information may be protected by security measures such as passwords and the like. Similarly, information stored on a non-volatile storage facility, such as an RFID disposed with the industrial machine, may include non-confidential information (e.g., serial number, model number and the like) that may be accessible to third-parties, and confidential

information (e.g., performance data, last failure date, prediction of next failure, failure rate of the machine or sup-portion thereof, and the like) that may require explicit authentication to access.

[1925] Accessing such rotating equipment information may include use of a mobile data collector, such as a mobile phone equipped with a data collection circuit that interacts with proximal industrial machines to access at least the non-confidential portion of the RFID tag. As the data collection circuit is activated to communicate with industrial machines, predictive maintenance beneficial information about the proximal industrial machines (e.g., as described herein and the like) may be collected from the RFID directly or by apply indexing (e.g., URL and the like) information gathered from the RFID to access the pertinent information from a networked server that is hosting the indexing information. In an example, a URL, which may be public data accessible in the RFID and a serial number of the machine, which may be treated as confidential information, may be retrieve from the RFID by the remote data collector. The data collector may provide the retrieved information to a predictive maintenance system that would apply the retrieved information in a web query to the URL, and the like.

[1926] Because some industrial machine deployments may not provide access to external networks like the Internet (e.g., for security purposes and the like), information in the RFID may be gathered and applied to predictive maintenance circuit operations contemporaneously with gathering the information; however predictive maintenance functions that require information not available at the time of gathering (e.g., information that must be retrieved over the Internet) may be performed at a later time, such as when the data collection circuit has access to the Internet and the like. In example embodiments, predictive maintenance event analysis may be performed on a suitably equipped data collection device (e.g., a mobile device with sufficient processing power and data storage, and the like) or on a server, such as a networked server and the like, or a combination thereof. Predictive maintenance event analysis may also be performed by computing equipment that is accessible over a network other than the Internet, such as a local area network that is accessible by the mobile data collector while in proximity to the industrial machine(s). Such a site-specific local area network may, with proper credentials presented from the mobile data collector, facilitate access to industrial machine rotating part-related information over the Internet and the like.

[1927] In example embodiments, rotor bar defects and weakening may be a precursor to secondary deterioration that can lead to further and costly repairs, such as replacement of a rotor core and the like. Therefore, by detecting broken or weakening rotor bars, maintenance and repair costs may be minimized. Knowing the count of rotor bars may be a factor in determining when maintenance and/or service of one or more rotor bars may be best actioned. As an example, by applying a rotor bar failure rate to a formula that predicts when a rotor bar may fail, knowing a count of rotor bars for a given machine, among other things like cycle rate, age, and the like can facilitate predicting when conducting service and/or testing of rotor bar-based systems could beneficially be conducted. A predictive maintenance circuit predicts maintenance events for industrial and other machines may predict maintenance for a machine with a greater number of rotor bars sooner than for a comparable machine with fewer rotor bars.

[1928] In example embodiments, predicting a maintenance event for a machine, such as a rotating equipment-based machine may be adapted from a predicted maintenance event for a similar machine while factoring in a count of gear teeth in the machine and the similar machine. An aspect of predicting the maintenance event that may be affected by, for example a count of gear teeth, may be a timing of the event. In an example, a machine with a greater number of gear teeth relative to the similar machine may suggest predicting a need for maintaining the machine with the greater number of gear teeth sooner than the similar machine. In example embodiments, predicting a maintenance event for a moving part of machine, such as a rotating equipment-based part may be adapted from a predicted maintenance event for a similar part in the same or similar machine while factoring in a count of gear teeth in the machine and the similar part or machine. In example embodiments, predicting a maintenance event for a rotating part of machine, such as a rotating part of a rotating equipment-based machine may be adapted from a predicted maintenance event for a similar rotating part in the same or similar machine while factoring in a count of gear teeth in the machine and the similar part or machine. In example embodiments, predicting a maintenance event for a gear box and the like, such as a rotating equipment-based gear box may be adapted from a predicted maintenance event for a similar part in the same or similar machine while factoring in a count of gear teeth in the machine and the similar part or machine. In example embodiments, predicting a maintenance event for a component of a machine comprising a multi-tooth gear, such as a rotating equipment-based component may be adapted from a predicted maintenance event for a similar component in the same or similar machine while factoring in a count of gear teeth in the machine and the similar component or machine.

[1929] In example embodiments, predicting a maintenance event for a rotating equipment may be a function of a predictive maintenance circuit that is, for example, responsive to a count of gear teeth of a rotatable component of a machine for which the predictive maintenance circuit products a maintenance event alert (e.g., a signal that facilitates triggering at least an automated portion of a maintenance event, such as ordering a replacement part and the like). In example embodiments, the predictive maintenance circuit may process operational data for the machine or rotating portion thereof, and/or may process failure data for a specific rotating component and the like of the machine or similar machines; thereby incorporating contextual information about the specific machine with static information about the machine such as gear teeth count and the like in the prediction.

[1930] In example embodiments, a count of gear teeth for a service component, such as from an RFID component integrated with or into an industrial machine, such as a rotary equipment, may be input to a machine learning circuit that may process the input along with service information for similar service components across a plurality of industrial machines. The machine learning circuit may generate a predictive maintenance adjustment factor that can be applied to the predictive maintenance circuit processing thereby producing a machine-specific predictive maintenance event.

[1931] In example embodiments, predicting a maintenance event for a rotating equipment may be a function of a predictive maintenance circuit that is, for example, responsive to a count of motor

rotor bars of a rotatable component of a machine for which the predictive maintenance circuit products a maintenance event alert. In example embodiments, a count of motor rotor bars for a service component, such as from an RFID component integrated with or into an industrial machine, such as a rotary equipment, may be input to a machine learning circuit that may process the input along with service information for similar service components across a plurality of industrial machines. The machine learning circuit may generate a predictive maintenance adjustment factor that can be applied to the predictive maintenance circuit processing thereby producing a machine-specific predictive maintenance event.

[1932] In example embodiments, predicting a maintenance event for a rotating equipment may be a function of a predictive maintenance circuit that is, for example, responsive to data representative of a revolutions per minute of, for example, an internal rotatable component of a machine for which the predictive maintenance circuit products a maintenance event alert. In example embodiments, RPM data for a service component, such as from an RFID component integrated with or into an industrial machine, such as a rotary equipment, may be input to a machine learning circuit that may process the input along with service information for similar service components across a plurality of industrial machines. The machine learning circuit may generate a predictive maintenance adjustment factor that can be applied to the predictive maintenance circuit processing thereby producing a machine-specific predictive maintenance event.

[1933] In example embodiments, predicting a maintenance event for a rotating equipment may be a function of a predictive maintenance circuit that is, for example, responsive to data representative of an aspect of a roller bearing, such as a number of balls per roller, a ball-to-roller contact angle, inner race dimensions, outer race dimensions, a number of vanes, a number of flutes, mode shape info, and the like of a rotatable component of a machine for which the predictive maintenance circuit products a maintenance event alert. In example embodiments, roller-bearing aspect data for a service component, such as from an RFID component integrated with or into an industrial machine, such as a rotary equipment, may be input to a machine learning circuit that may process the input along with service information for similar service components across a plurality of industrial machines. The machine learning circuit may generate a predictive maintenance adjustment factor that can be applied to the predictive maintenance circuit processing thereby producing a machine-specific predictive maintenance event. In example embodiments, a predicted maintenance event may be selected from a list of maintenance events including, without limitation part replacement, machine sub-system replacement, calibration, deep data collection, machine servicing, machine shutdown, preventive maintenance, and the like.

[1934] In example embodiments, at least one aspect of a roller bearing service component may be stored in a portion of digital data structure of roller bearing component production information retrieved through an RFID component disposed with the roller bearing component into an industrial machine. In example embodiments, the portion of the digital data structure may be specific to the industrial machine with which the roller bearing component is disposed. In example embodiments, the portion of the digital data structure may be retrieved by accessing a network location retrieved from the RFID component and further indexed by a machine-specific identifier

retrieved from the RFID component. In example embodiments, the network location may be accessed through a Wi-Fi interface of a data collection device while the data collection device is in short range wireless communication with the RFID component. Further In example embodiments, the network location may be accessed through a Wi-Fi interface of a data collection device when the data collection device is no longer in short range wireless communication with the RFID component. In example embodiments, the portion of the digital data structure may be retrieved by providing a machine-specific key retrieved from the RFID component to an Application Programming Interface function of a predictive maintenance system that facilitates access to roller bearing component production information stored external to the industrial machine. In example embodiments, the portion of the digital data structure may include production information retrieved from the RFID component. In example embodiments, the circuit predicts a maintenance event for the roller bearing component responsive to retrieving the portion of the digital data structure from the RFID component independent of network connectivity of a processor executing the circuit. Yet further in example embodiments, a data collection device may include the predictive maintenance circuit that predicts a maintenance event for the roller bearing component responsive to retrieving the portion of the digital data structure from the RFID component independent of network connectivity of the data collection device.

[1935] Referring to Figure 201, a diagram of a data structure 31000 for storing rotating part-related information for use in, among other things, predicting a maintenance event for a portion of an industrial machine associated with the rotating part is depicted. A rotating component 31002 may include a specific gear of an industrial machine, a gear in a gearbox, a shaft, roller bearings and the like. Parameters 31004 for each rotating component may include, without limitation, count of teeth, count of gears, type(s) of gears in a gear box, rotation rate, count of balls, race dimensions, number of vanes and the like. Values 31006 for each rotating component-parameter combination may be stored in the data structure 31000. This data structure may be representative of a portion of rotating part data stored on an RFID component deployed with an industrial machine. The number of entries on the data structure, types of data in the data structure, and formats for values (e.g., decimal, hexadecimal, and the like) may vary as needed to support storing rotating part-related configuration, production and test information.

[1936] Referring to Figure 202, a flow chart is depicted that represents a method for predicting a maintenance event for a rotating part, such as a gear, motor, roller bearing and the like based on as stream of sensed rotating part health data and part-specific configuration information, such as gear tooth count, roller bearing/chase dimensions, rotor bar count for a motor, and the like. A method 31100 may include a step 31102 of generating streams of health data for a rotating part, such as a gear, motor, roller bearing and the like. The method 31100 may continue with a step 31104 of accessing configuration information for the rotating part, such as from an RFID part deployed with the industrial machine hosting the rotating part and/or from a network-accessible data storage facility. The method 31100 may continue with a step 31106 of predicting at least one of a gear, motor, and/or roller bearing related maintenance event/action/likelihood. The method 31100 may continue with a step 31108 of producing orders for the predicted maintenance action to maintain,

repair, and/or replace the rotating part for which a maintenance action/event is predicted. The method 31100 may continue with a step 31110 of validating the maintenance action(s) taken based on the rotating part based on service data for the maintenance event; such data for the maintenance event may be received by a processor, such as a networked server from the industrial machine and the like.

[1937] The present disclosure is also related to an Industrial Internet of Things (IIoT) system that is configured to address the above identified and other needs. More particularly, the present disclosure is directed to an IIoT platform that is optimized to improve the collection, storage, processing, sharing, and utilization of data in an industrial environment. The IIoT platform can be arranged in a plurality of distinct data-handling layers in a layered topology. This layered topology facilitates independent optimization of each of the data-handling layers. For example only, the layers can include a data collection/monitoring layer, a data storage layer, an adaptive intelligence layer, and an application platform layer. Each of the layers can have a micro-services architecture and interfaces to the other layers such that outputs, events, outcomes, etc. can be exchanged and shared across the layers. In this manner, and as mentioned above, each of the data-handling layers can be independently optimized for their specific functions (storage, monitoring, intelligence development, and applications) while permitting cross-layer sharing and optimization of the platform as a whole.

[1938] In one aspect, the IIoT platform can comprise a multi-application IIoT application platform that shares a common infrastructure that facilitates intelligence development and utilization. The common infrastructure provides for cross-application and cross-layer data sharing, including the sharing of events, outputs, and outcomes, to facilitate coordinated optimization (e.g., via machine learning) of the IIoT platform. The common data handling infrastructure can enable efficient monitoring of industrial entities and applications, as well as efficient sharing of such gathered data, to provide an environment for rapid development and deployment of intelligence solutions. The common infrastructure can also provide a consistent user experience for multiple applications related to different industrial processes.

[1939] In another aspect, the IIoT platform can include an adaptive intelligence layer that provides adaptive intelligence solutions to the various components in the IIoT platform. The adaptive intelligence layer can include a set of data processing, artificial intelligence, and computational systems that develop, improve, or adapt processes in the IIoT platform. The adaptive intelligence layer utilizes data collected, generated, stored, or otherwise obtained by the IIoT platform. The data can, for example, be related to various entities in the industrial environment, including but not limited to machines, devices, processes, workflows, and combinations thereof. The adaptive intelligence layer can include an adaptive edge compute management system that adaptively manages edge computation, storage, and processing in the IIoT system. Additionally or alternatively, the adaptive intelligence layer can include a robotic process automation system that develops and deploys automation capabilities for at least one of the plurality of industrial entities in the IIoT system. Further, the adaptive intelligence layer can include a set of protocol adaptors that facilitate adaptive protocol transformations of data within the IIoT system. The adaptive

intelligence layer can additionally or alternatively include an edge intelligence system that adapts edge computation resources. For example only, the edge intelligence system can adapt the edge computation resources such that computational resources are utilized in an optimized manner based on various constraints (speed, cost, etc.).

[1940] The adaptive intelligence layer can, in further aspects, include an adaptive networking system that adapts network communication in the IIoT system. In other aspects, the adaptive intelligence layer can include a set of state and event managers that adapt the processes in the IIoT system based on state and event data. An opportunity mining system (which may include and also be referred to herein as a set of opportunity miners) can also be included in the adaptive intelligence layer. The set of opportunity miners can identify opportunities for increased automation or intelligence in the IIoT system. Finally, the adaptive intelligence layer can include a set of artificial intelligence systems that develop, improve, or adapt processes in the IIoT system.

[1941] As mentioned above, the robotic process automation system develops and deploys automation capabilities for at least one of the plurality of industrial entities in the IIoT system. The robotic process automation system can develop such capabilities for each of the processes, workflows, etc. that is managed, controlled, or mediated by each of the applications in the multi-application IIoT application platform. Further, the robotic process automation system can develop such capabilities for combinations of the applications. Additionally or alternatively, the robotic process automation system can develop and deploy automation capabilities for various industrial processes, including but not limited to energy production processes, manufacturing processes, transport processes, storage processes, refining processes, distilling processes, fluid handling processes, energy storage processes, chemical processes, petrochemical processes, semiconductor processes, gas production processes, maintenance processes, service processes, repair processes, and supply chain processes.

[1942] The robotic process automation system can develop and deploy automation capabilities based on watching/monitoring software interactions (e.g., by workers with various software interfaces), hardware interactions (e.g., by watching workers actually interacting with or using machines, equipment, tools or the like), or combinations thereof. Further, the robotic process automation system can utilize data gathered, generated, or otherwise obtained from or about the IIoT platform to assist in its activities.

[1943] As briefly mentioned above, the set of protocol adaptors facilitate adaptive protocol transformations of data within the IIoT system. For example only, the set of protocol adaptors can facilitate adaptive in-flight data protocol transformations, communication network protocol transformations, and linking (gateways, routers, switches, etc.). In some aspects, this includes recognition of appropriate protocols used by various components and systems in each of the data handling layers and in each industrial environment such that data can be moved, stored, and processed regardless of the native storage format, processing format, or communication system protocol. In some aspects, the set of protocol adaptors can be self-organizing. The self-organizing protocol adaptor can facilitate adaptive in-flight data protocol transformation of the data by selecting at least one interface of a set of possible interfaces between communication nodes.

Alternatively or additionally, the self-organizing protocol adaptor can facilitate adaptive in-flight data protocol transformation of the data by selecting an appropriate protocol for the data and, in some aspects, also transform the data to comply with the selected appropriate protocol.

[1944] As mentioned above, the adaptive intelligent systems layer can include an opportunity mining system that utilizes the data to identify opportunities for increased automation within the platform. The opportunity mining system can be configured to collect information within the platform and also within, about, and for a set of industrial environments and industrial entities that help identify and prioritize opportunities for increased automation and/or intelligence in the IIoT system. The opportunity mining system can, for example, utilize sensors (such as cameras or wearables) or other systems to observe clusters of workers by time, by type, and by location to identify labor-intensive areas and processes. Further, the opportunity mining system can characterize the extent of domain-specific or entity-specific knowledge or expertise required to undertake an action, use a program, use a machine, or the like, such as observing the identity, credentials, and experience of workers involved in given processes. Alternatively or additionally, in some implementations the opportunity mining system can include systems by which a developer can solicit or specify information that would be helpful (such as video showing an expert doing something) and provide consideration/rewards for providing the specified information.

[1945] In certain aspects, the adaptive intelligent systems layer can include an edge intelligence system that adapts edge computation resources. The edge intelligence system can adaptively manage "edge" computation, storage, and processing, such as by varying storage locations for data and processing locations (e.g., applying AI) between on-device storage, local systems, in the network, and in the cloud. The edge intelligence system can permit and facilitate the dynamic definition of what constitutes the "edge" for purposes of a given application, device, system, etc. Further, the edge intelligence system can permit adaptation of edge computation that is multi-application aware, such as accounting for Quality of Service, latency requirements, congestion, cost, and other factors.

[1946] In other aspects, the industrial entity-oriented data storage systems layer can include at least one geofenced virtual asset tag associated with one particular industrial entity of the plurality of industrial entities in the IIoT system. The at least one geofenced virtual asset tag can comprise a data structure that contains entity data about the one particular industrial entity and is linked to the proximity of the one particular industrial entity. Essentially, a geofenced virtual asset tag limits access as if the tag were physically located on an asset. IIoT devices within the geofence can be used to recognize the presence of a reader device (such as by recognition of an interrogation signal) and communicate, e.g., with help of protocol adaptors, with the geofenced virtual asset tag. Further, in some aspects, IIoT devices can act as distributed blockchain nodes, such as for validation (such as by various consensus protocols) of enchained data, including transaction history for maintenance, repair, and service. IIoT devices in the geofence can collectively validate location and identity of a fixed asset, e.g., in a configuration in which neighbors validate other neighbors.

[1947] Referring to Figure 203, a platform 34900 for facilitating development of intelligence in an Industrial Internet of Things (IIoT) system is illustrated, including a set of systems, applications,

processes, modules, services, layers, devices, components, machines, products, sub-systems, interfaces, connections, and other elements working in coordination to enable intelligent management of a set of industrial entities 34930 that may be part of, integrated with, linked to, or operated on by the platform 34900. Industrial entities 34930 may include any of the wide variety of assets, systems, devices, machines, facilities, individuals, or other entities mentioned throughout this disclosure or in the documents incorporated herein by reference, such as, without limitation: industrial machines 34952 and their components (factory components, power production machinery, turbines, motors, reactors, fluid handling systems, condensers, fans, software components, hardware components, electrical components, physical components, etc.); industrial processes 34950 (power production processes, software processes (including applications, programs, services, and others), factory production processes, manufacturing processes (e.g., semiconductor manufacturing processes, chemical manufacturing processes, petroleum manufacturing processes, biological manufacturing processes), service, maintenance and repair processes, diagnostic processes, security processes, safety processes and many others); wearable and portable devices 34948 (mobile phones, tablets, dedicated portable devices for industrial applications, data collectors (including mobile data collectors), sensor-based devices, watches, glasses, hearables, head-worn devices, clothing-integrated devices, arm bands, bracelets, neck-worn devices, AR/VR devices, headphones, etc.); workers 34944 (factory workers, maintenance and service personnel, managers, engineers, floor managers, warehouse workers, inspectors, refueling personnel, material handling workers, process supervisors, security personnel, safety personnel, etc.); robotic systems 34942 (physical robots, collaborative robots ("cobots"), software bots, etc.); and operating facilities 34940 (power production facilities, refineries, assembly facilities, manufacturing facilities, warehousing facilities, plants, factories, mining facilities, power extraction facilities, construction sites, exploration sites, drilling sites, harvesting sites, etc.), which may include, without limitation, storage and warehousing facilities (such as for warehousing inventory, components, packaging materials, goods, products, machinery, equipment, and other items); transportation facilities 34934 (ports, depots, hangars, transportation equipment, vehicles, docks, loading bays, assembly lines, and other facilities for moving goods, components, machinery, raw materials, and other items); and manufacturing facilities 34932 (such as for manufacturing, assembling, refining, finishing, packaging, or otherwise producing a wide variety of goods).

[1948] In example embodiments, the platform 34900 may include a plurality of data handling layers 34908, each of which being configured to provide a set of capabilities that facilitate development and deployment of intelligence (such as for facilitating automation, machine learning, applications of artificial intelligence, intelligent transactions, state management, event management, and process management) for a wide variety of industrial applications and end uses. In some implementations, the data handling layers 34908 include an industrial monitoring systems layer 34906, an industrial entity-oriented data storage systems layer 34910 (referred to in some cases herein for convenience simply as a data storage layer 34910), an adaptive intelligent systems layer 34904, and an industrial management application platform layer 34902. Each of the data

handling layers 34908 may include a variety of services, programs, applications, workflows, systems, components and modules, as further described herein and in the documents incorporated herein by reference. In certain implementations, each of the data handling layers 34908 (and optionally the platform 34900 as a whole) is configured such that one more of its elements can be accessed as a service by other layers 34908 or by other systems, e.g., by being configured as a platform-as-a-service deployed on a set of cloud infrastructure components in a microservices architecture. For example only, a data handling layer 34908 may have a set of interfaces 34980 (application programming interfaces (APIs), brokers, services, connectors, wired or wireless communication links, ports, human-accessible interfaces, software interfaces or the like) by which data may be exchanged between the data handling layer 34908 and other layers, systems or sub-systems of the platform 34900, as well as with other systems (such as industrial entities 34930 or external systems, cloud-based or on-premises enterprise systems (e.g., accounting systems, resource management systems, customer-relationship management (CRM) systems, and supply chain management systems). Each of the data handling layers 34908 may include a set of services (e.g., microservices) for data handling, including facilities for data extraction, transformation, and loading; data cleansing and deduplication facilities; data normalization facilities; data synchronization facilities; data security facilities; computational facilities (e.g., for performing pre-defined calculation operations on data streams and providing an output stream); compression and de-compression facilities; and analytic facilities (such as providing automated production of data visualizations).

[1949] In various aspects, each data handling layer 34908 has a set of interfaces 34980 (such as application programming interfaces or "APIs") for automating data exchange with each of the other data handling layers 34908. In aspects, the data handling layers 34908 are configured in a topology that facilitates shared data collection and distribution across multiple applications and uses within the platform 34900 by the industrial monitoring systems layer 34906. The industrial monitoring systems layer 34906 may include various data collection and management systems 34918 (referred to for convenience in some cases as data collection systems 34918) for collecting and organizing data collected from or about industrial entities 34930, as well as data collected from or about the various data layers 34908 or services and/or components thereof.

[1950] For example, a stream of physiological data from a wearable device worn by a worker 34944 on a factory floor can be distributed via the industrial monitoring systems layer 34906 to multiple distinct applications in the industrial management application platform layer 34902, such as one that facilitates monitoring the health of a worker and another that facilitates operational efficiency. In aspects, the industrial monitoring systems layer 34906 facilitates alignment (such as time-synchronization, normalization, or the like) of data that is collected with respect to one or more industrial entities 34930. For example, one or more video streams collected of a worker 34944 in an industrial environment, such as from a set of camera-enabled IoT devices, may be aligned with a common clock, so that the relative timing of a set of videos can be understood by systems that may process the videos, such as machine learning systems that operate on images in the videos, on changes between images in different frames of the video, or the like. In such an

example, the industrial monitoring systems layer 34906 may further align a set of videos with other data, such as a stream of data from wearable devices, a stream of data produced by industrial systems (such as on-board diagnostic systems, telematics systems, and various other sensors), a stream of data collected by mobile data collectors, and any other data or data stream sensed, generated, or otherwise obtained. Configuring the industrial monitoring systems layer 34906 as a common platform (or set of microservices) that are accessed across many applications may dramatically reduce the number of interconnections required by an enterprise in order to have a growing set of applications monitoring a growing set of IoT devices and other systems and devices that are under its control.

[1951] In aspects, the data handling layers 34908 are configured in a topology that facilitates shared or common data storage across multiple applications and uses of the platform 34900 by the industrial entity-oriented data storage systems layer 34910, referred to herein for convenience in some cases simply as the storage layer 34910. For example, various data collected about the industrial entities 34930, as well as data produced by the other data handling layers 34908, may be stored in the industrial entity-oriented data storage systems layer 34910, such that any of the services, applications, programs, etc. of the various data handling layers 34908 can access a common data source. This may facilitate a dramatic reduction in the amount of data storage required to handle the enormous amount of data produced by or about industrial entities 34930 in the platform 34900. For example, a supply chain management application in the industrial management application platform layer 34902 (such as one for ordering replacement parts) may access the same data set about what parts have been replaced for a set of machines as a predictive maintenance application that is used to predict whether a machine is likely to require repairs. In aspects, the industrial entity-oriented data storage systems layer 34910 may provide an extremely rich environment for collection of data that can be used for extraction of features or inputs for intelligence systems, such as expert systems, artificial intelligence systems, robotic process automation systems, machine learning systems, deep learning systems, supervised learning systems, or other intelligent systems as disclosed throughout this disclosure and the documents incorporated herein by reference. As a result, each application in the industrial management application platform layer 34902 and each adaptive intelligent system in the adaptive intelligent systems layer 34904 can benefit from the data collected or produced by or for each of the others.

[1952] A wide range of data types may be stored in the storage layer 34910 using various storage media and data storage types and formats, including, without limitation: asset and facility data 34920 (including asset identity data, operational data, transactional data, event data, state data, workflow data, maintenance data, and other data); worker data 34922 (including identity data, role data, task data, workflow data, health data, performance data, quality data, and other data); event data 34924 (including data regarding process events, financial events, output events, input events, state-change events, operating events, repair events, maintenance events, service events, damage events, injury events, replacement events, refueling events, recharging events, supply events, and others); claims data 34954 (including data related to insurance claims, such as for business interruption insurance, product liability insurance, insurance on goods, facilities, or equipment,

flood insurance, insurance for contract-related risks, and others; data related to product liability, general liability, workers compensation, injury, and other liability claims; and claims data relating to contracts, such as supply contract performance claims, product delivery requirements, warranty claims, indemnification claims, energy production requirements, delivery requirements, timing requirements, milestones, key performance indicators, and others); production data 34958 (such as data relating to energy production found in databases of public utilities or independent services organizations that maintain energy infrastructure; data relating to outputs of manufacturing; data related to outputs of mining and energy extraction facilities, drilling and pipeline facilities, and many others); and supply chain data 34960 (such as data related to items supplied, amounts, pricing, delivery, sources, routes, customs information, and other supply chain facets).

[1953] In aspects, the data handling layers 34908 are configured in a topology that facilitates shared adaptation capabilities, which may be provided, managed, mediated, etc. by one or more of a set of services, components, programs, systems, or capabilities of the adaptive intelligent systems layer 34904, referred to in some cases herein for convenience as the adaptive intelligence layer 34904. The adaptive intelligence systems layer 34904 may include a set of data processing, artificial intelligence, and computational systems 34914 that are described in more detail elsewhere throughout this disclosure. Thus, use of various resources, such as computing resources (available processing cores, available servers, available edge computing resources, available on-device resources - for single devices or peered networks, available cloud infrastructure, etc.), data storage resources (including local storage on devices, storage resources in or on industrial entities or environments (including on-device storage, storage on asset tags, local area network storage), network storage resources, cloud-based storage resources, database resources, and others), networking resources (including cellular network spectrum, wireless network resources, fixed network resources, and others), energy resources (available battery power, available renewable energy, fuel, grid-based power, etc.), may be optimized in a coordinated or shared way on behalf of an operator, enterprise, system, application, or the like, such as for the benefit of multiple applications, programs, workflows, or other services/processes. For example, the adaptive intelligence layer 34904 may manage and provision available network resources for both an industrial analytics application and for an industrial remote control application such that low latency resources are used for remote control and longer latency resources are used for the analytics application. As described in more detail throughout this disclosure and the documents incorporated herein by reference, a wide variety of adaptations may be provided on behalf of the various services and capabilities across the various layers 34908, including ones based on application requirements, quality of service, budgets, costs, pricing, risk factors, operational objectives, optimization parameters, returns on investment, profitability, and uptime/downtime.

[1954] The industrial management application platform layer 34902, referred to in some cases herein for convenience as the application platform layer 34902, may include a set of industrial processes, workflows, activities, events, and applications 34912 (referred to individually and collectively, except where context indicates otherwise, as applications 34912) that enable an operator to manage more than one aspect of an industrial environment or industrial entity 34930 in

a common application environment. The common application environment may permit the platform 34900 to take advantage of common data storage in the data storage layer 34910, common data collection or monitoring in the industrial monitoring systems layer 34906, and/or common adaptive intelligence of the adaptive intelligence systems layer 34904. Outputs from the applications 34912 in the application platform layer 34902 may be provided to the other data handing layers 34908. These may include, without limitation, state and status information for various objects, entities, processes, flows and the like; object information (such as identity, attribute, and parameter information for various classes of objects of various data types); event and change information (such as for workflows, dynamic systems, processes, procedures, protocols, and algorithms) including but not limited to timing information; outcome information (such as indications of success and failure, indications of process or milestone completion, indications of correct or incorrect predictions, indications of correct or incorrect labeling or classification, and success metrics such as those relating to yield, engagement, return on investment, profitability, efficiency, timeliness, quality of service, quality of product, customer satisfaction, and other measures of success). Outputs from each application 34912 can be stored in the data storage layer 34910, distributed for processing by the data collection layer 34906, and/or used by the adaptive intelligence layer 34904. The cross-application nature of the application platform layer 34902 thus facilitates convenient organization of all of the necessary infrastructure elements for adding intelligence to any given application, such as by supplying machine learning on outcomes across applications, providing enrichment of automation of a given application via machine learning based on outcomes from other applications (or other elements of the platform 34900), and allowing application developers to focus on application-native processes while benefiting from other capabilities of the platform 34900.

[1955] Referring to Figure 204, additional details, components, sub-systems, and other elements of an optional implementation of the platform 34900 of Figure 203 are illustrated. The industrial management application platform layer 34902 can include, in various optional implementations, a set of applications, systems, solutions, interfaces, or services (for convenience, referred to herein individually and collectively as applications 34912), by which an operator or owner of an industrial entity 34930, or other user, may manage, monitor, control, analyze, or otherwise interact with one or more elements of the industrial entity 34930. The set of applications 34912 may include one or more other applications 34912 that facilitates improved operation of an industrial entity, facility, or the like for the owner, operator, or other user, including but not limited to one or more of a blockchain-based industrial asset lifecycle management application 35002, an industrial asset lifecycle management application 35004, a process control optimization application 35010, a building automation and controls application 35012, an enterprise asset management application 35014, a cloud/PaaS/SaaS solution 35008, a factory operations visibility and intelligence (FOVI) application 35018, an autonomous manufacturing application 35020, a smart supply chain application 35022, an inventory quality control application 35024, and an industrial analytics application 35028.

[1956] In certain aspects, the one or more applications 34912 of the industrial management application platform layer 34902 and/or the artificial intelligence systems 35048 can include an artificial intelligence-enabled assistant 35089 that provides documentation related to an industrial entity 34930 (such as a machine and/or process that may require maintenance or repair), that provides diagnostics on the industrial entity 34930, and/or provides a set of recommendations for service, update, maintenance, replacement, repair, or other activity. This artificial intelligence-enabled assistant 35089 can be part of a suite of solutions or applications 34912 that use capabilities of the platform 34900 and the various shared microservices and layers (including artificial intelligence and advanced analytics) to enable preventative and predictive tasks related to the industrial entity 34930, such as downtime and maintenance management.

[1957] In further aspects, the applications 34912 can also include an asset performance management solution 35091 and/or an enterprise asset management application 35093 to, among other things, reduce the risk of failure or improve performance of various assets or industrial entities 34930, such as vehicles, manufacturing robots, turbines, mining equipment, elevators, transformers, motors, generators, and other machines or components thereof. Such solutions can use the data collection systems 34918 and other data sources to collect data from physical assets in near real-time and to provide information regarding operating conditions, process status, and/or fault conditions, as well as predict potential issues and other similar tasks. In aspects, recommendations can be provided for service, maintenance, repair, updates, or replacement, including, as described throughout this disclosure and the documents incorporated by reference, recommendations as to replacement parts, procedural information, identification of timing and schedule information, identification of personnel or entities capable of undertaking repairs, ratings, and other similar information.

[1958] In various implementations, applications 34912 may include industry-specific or entity-specific versions, such as for the energy industry, manufacturing industries, power generation industries, and mining industries. It should be appreciated that other entities/industries are contemplated and fall within the scope of the present disclosure. The data collected, organized, compiled, generated, utilized, etc. by the industry-specific or entity-specific versions can include industry specific risk models, models for performance and degradation of particular types of machines, and external data, such as on weather conditions, operational conditions, and/or market conditions.

[1959] In some implementations, hardware for machine learning at the edge can take the form of a single-board computer running an edge-based Tensor Processing Unit (TPU), as well as a system-on-module (SOM) (such as the recently announced SOM available from Coral™), and/or a USB-connected or other accessory device that brings machine learning inferencing to existing systems.

[1960] In certain aspects, the adaptive intelligent systems layer 34904 may include a set of systems, components, services, and other capabilities that collectively facilitate the coordinated development and deployment of intelligent systems, such as ones that can enhance one or more of the applications 34912 at the industrial management application platform layer 34902. The adaptive intelligence systems layer 34904 can include, for example, an adaptive edge compute

management system 35030, a robotic process automation system 35042, a set of protocol adaptors 35602, a packet acceleration system 35034, an edge intelligence system 35038, an adaptive networking system 35040, a set of state and event managers 35044, a set of opportunity miners 35046, and a set of artificial intelligence systems 35048, although additional or fewer elements are possible.

[1961] In aspects, the industrial monitoring systems layer 34906 and its data collection systems 34918 may include a wide range of systems for collection of data. This layer may include, without limitation, real time monitoring systems 35068 (such as onboard monitoring systems like on-board diagnostics and telematics systems, monitoring infrastructure (such as cameras, motion sensors, and ambient sensors), as well as removable and replaceable monitoring systems, such as portable and mobile data collectors); software interaction observation systems 35050 (such as for logging and tracking events involved in interactions of users with software user interfaces (mouse movements, mouse clicks, cursor movements, keyboard interactions, navigation actions, eye movements, menu selections, etc.), as well as software interactions that occur as a result of other programs, such as over APIs); mobile data collectors 35052 (such as described herein and in documents incorporated by reference), visual quality detection systems 35054 (including use of video and still imaging systems, LIDAR, IR and other systems that allow visualization of materials, components, machines, housings, seals, bearings, and many other elements of industrial entities 34930, as well as inspection systems that monitor processes, activities of workers, and the like); on-board diagnostic (OBD) and telematics systems 35070 that can provide diagnostic codes and events via an event bus, communication port, or other communication system; physical process observation systems 35058 such as for tracking physical interactions of workers with other workers, workers with physical entities like machines and equipment, and physical entities with other physical entities, including, without limitation, video cameras, motion sensing systems (such as including optical sensors, LIDAR, IR and other sensor sets), and robotic motion tracking systems (such as tracking movements of systems attached to a human or a physical entity); machine condition monitoring systems 35060 (including onboard monitors and external monitors of conditions, states, operating parameters, or other measures of the condition of a machine); sensors and cameras 35062 (including onboard sensors, sensors in an industrial environment, cameras for monitoring an entire environment, dedicated cameras for a particular machine, process, worker, or other feature, wearable cameras, portable cameras, cameras disposed on mobile robots, cameras of portable devices like smart phones and tablets, and any of the many sensor types disclosed throughout this disclosure or in the documents incorporated herein by reference); indoor air quality monitoring systems 35072 (including chemical noses and other chemical sensor sets, as well as visual sensors); continuous emission monitoring systems 35074; indoor sound monitoring systems 35078; and any other of a wide variety of Internet of Things (IoT) data collectors, such as those described throughout this disclosure and in the documents incorporated by reference herein.

[1962] In certain implementations, and as mentioned above, the industrial entity-oriented data storage systems layer 34910 can include a range of systems for storage of data. These may include, without limitation, physical storage systems, virtual storage systems, local storage systems 35092,

distributed storage systems, databases, memory, network-based storage, and network-attached storage systems 35082 (such as using non-volatile memory express ("NVMe"), storage attached networks, and other network storage systems). Additionally or alternatively, the storage layer 34910 may store data in one or more knowledge graphs 35080, such as a directed acyclic graph, a data map, a data hierarchy, or a self-organizing map. Further, the data storage layer 34910 may store data in an industrial digital thread 35084, such as for maintaining a longitudinal record of an industrial entity 34930 over time, including any of the entities described herein. As described further herein, the data storage layer 34910 may use and enable a virtual asset tag 35088, which may include a data structure that is associated with an asset and accessible and managed as if the tag were physically located on the asset, such as by use of access controls, so that storage and retrieval of data is optionally linked to local processes, but also optionally open to remote retrieval and storage options. In embodiments the storage layer 34910 may include one or more blockchains 35090, such as ones that store identity data, transaction data, historical interaction data, and other data, such as with access control that may be role-based or may be based on credentials associated with an industrial entity 34930, a service, or one or more applications 34912.

[1963] With further reference to Figure 205, the adaptive intelligence systems layer 34904 may include a robotic process automation ("RPA") system 35042 that includes a set of components, processes, services, interfaces, and other elements for development and deployment of automation capabilities for various industrial entities 34930, environments, and applications 34912. Without limitation, the robotic process automation system 35042 may apply automation capabilities to each of the processes that is managed, controlled, or mediated by each of the set of applications 34912 of the application platform layer 34902.

[1964] In aspects, the robotic process automation system 35042 may leverage the presence of multiple applications 34912 within the industrial management application platform layer 34902 such that a pair of applications may share data sources (such as in the data storage layer 34910) and other inputs (such as from the industrial monitoring systems layer 34906) that are collected with respect to industrial entities 34930, as well sharing outputs (such as events, state information, and other data), which collectively may provide a much richer environment for process automation, including through the use of artificial intelligence systems 35048 (including any of the various expert systems, artificial intelligence systems, neural networks, supervised learning systems, machine learning systems, deep learning systems, and other systems described throughout this disclosure and in the documents incorporated by reference).

[1965] For example, an inventory quality control application 35024 may use the robotic process automation system 35042 for automation of an inspection process that is normally performed or supervised by a human. The process could involve visual inspection using video or still images from a camera or other imaging device that displays images of an entity 34930, such as where the robotic process automation 35042 system is trained to automate the inspection by observing interactions of a set of human inspectors or supervisors with an interface that is used to identify, diagnose, measure, parameterize, or otherwise characterize possible defects in an item. In aspects, the interactions of the human inspectors or supervisors may include a labeled data set where labels

or tags indicate types of defects or other characteristics such that a machine learning system can learn, using the training data set, to identify the same characteristics. The identification of the same characteristics can, in turn, be used to automate the visual quality detection process such that defects are automatically classified and detected in a set of video or still images, which in turn can be used within the inventory quality control application 35024 to flag items of inventory that should be rejected or otherwise require further inspection. In certain implementations, the robotic process automation system 35042 may involve multi-application or cross-application sharing of inputs, data structures, data sources, events, states, outputs, or outcomes. For example, the inventory quality application 35042 may receive information from a smart supply chain application 35022 in order to enrich the robotic process automation by the robotic process automation system 35042 of the inventory quality control application 35042, such as information about the expected characteristics of a product or other item from a particular vendor, which may assist in reducing false positive or false negatives in a visual inspection process. These and many other examples of multi-application or cross-application sharing for robotic process automation 35042 across the applications 34912 are encompassed by the present disclosure.

[1966] In various implementations, the robotic process automation system 35042 may operate on shared or converged processes among the various pairs of the applications 34912 of the industrial management application platform layer 34902, such as, without limitation, of a converged process involving factory operations visual intelligence (FOVI) system 35018 and process control optimization (PCO) system 35010, and integrated automation of blockchain-based industrial asset lifecycle management application 35002 with smart supply chain application 35022. Other examples are contemplated by this disclosure.

[1967] In certain aspects, the converged processes may include shared data structures for multiple applications 34912, including ones that track the same transactions on a blockchain but may consume different subsets of available attributes of the data objects maintained in the blockchain or ones that use a set of nodes and links in a common knowledge graph. For example, a transaction indicating a change of ownership of an industrial entity 34930 may be stored in a blockchain and used by multiple applications 34912, such as to enable role-based access control, role-based permissions for remote control, identity-based event reporting, and other functions. In aspects, converged processes may include shared process flows across applications 34912, including subsets of larger flows that are involved in one or more of a set of applications 34912. For example, a visual inspection flow about an entity 34930 may serve an inventory quality control application 35024, an industrial analytics application 35028, an enterprise asset management application 35014, and others.

[1968] In example embodiments, the RPA system 35042 may provide robotic process automation for the wide range of industrial processes mentioned throughout this disclosure and the documents incorporated herein by reference, including without limitation energy production, manufacturing, transport, storage, refining, distilling, fluid handling, energy storage, chemical processes, petrochemical processes, semiconductor processes, gas production processes, maintenance processes, service processes, repair processes, supply chain processes, assembly line processes,

inspection processes, purchase and sale processes, fault detection processes, and power utilization optimization processes.

[1969] An environment for development of robotic process automation may include a set of interfaces for developers in which a developer may configure an artificial intelligence system 35048 to take inputs from selected data sources of the data storage layer 34910 and events or other data from the industrial monitoring systems layer 34906 and supply them, such as to a neural network, either as inputs for classification or prediction, as outcomes, or for other purposes to the RPA system 35042. The RPA system 35042 may be configured to take one or more process and application outputs and outcomes 34928 from various applications 34912 to facilitate automated learning and improvement of classification, prediction, or other activities that are involved in a process that is intended to be automated.

[1970] In aspects, the development environment, and the resulting robotic process automation performed by the RPA system 35042, may involve monitoring a combination of both software program interaction observations (e.g., received from the software interaction observation systems 35050), such as by observing workers interacting with various software interfaces of applications 34912 involving industrial entities 34930, and physical process interaction observations (e.g., received from the physical process observation systems 35058), such as by watching workers interacting with or using machines, equipment, tools, or other components. In various implementations, observation of software interactions by the software interaction observation systems 35050 may include observation of interactions among software components with other software components, such as how one application 34912 interacts via APIs with another application 34912. In certain aspects, observation of physical process interactions by the physical process observation systems 35058 may include observation (such as by video cameras, motion detectors, or other sensors) as well as detection of various physical interactions between industrial entities 34930 and/or its individual elements. For example only, such physical interactions can include without limitation observation/detection of positions, movements, and the like of hardware (such as robotic hardware), how human workers interact with industrial entities 34930 (such as locations of workers, including routes taken through a facility, where workers of a given type are located during a given set of events, processes or the like, how workers manipulate pieces of equipment or other items using various tools and physical interfaces, the timing of worker responses with respect to various events (e.g., responses to alerts and warnings), procedures by which workers undertake scheduled maintenance, updates, repairs, and service processes, procedures by which workers tune or adjust items involved in production). Physical process observation systems 35058 may track positions, angles, forces, velocities, acceleration, pressures, torque, and other characteristics of a worker as the worker operates on hardware (such as with a tool). Such observations may be obtained by any combination of video data, data detected within a machine (such as of positions of elements of the machine detected and reported by position detectors), data collected by a wearable device (such as an exoskeleton that contains position detectors, force detectors, torque detectors, and/or other sensors that is configured to detect the physical characteristics of interactions of a human worker with a hardware item for purposes of

developing a training data set). By collecting both software interaction observations (e.g., with software interaction observation systems 35050) and physical process interaction observations (e.g., with physical process observation systems 35058), the RPA system 35042 can more comprehensively automate processes involving industrial entities 34930, such as by using software automation in combination with physical robots.

[1971] In various implementations, the RPA system 35042 is configured to train a set of physical robots that have hardware elements that facilitate undertaking tasks that are conventionally performed by humans. These may include robots that, among other activities, walk (including walking up and down stairs), climb (such as climbing ladders), move about a facility, attach to items, grip items (such as using robotic arms, hands, pincers, or the like), lift items, carry items, remove and replace items, and use tools.

[1972] Referring to Figure 206, an opportunity mining system 35046 may be provided as part of the adaptive intelligence layer 34904. The opportunity mining system 35046 may be configured to seek and recommend opportunities to improve one or more of the elements of the platform 34900, such as via addition of artificial intelligence systems 35048, automation (including robotic process automation, e.g., via robotic process automation system 35046 or otherwise), or the like to one or more of the systems, sub-systems, components, applications, or other aspects of the platform 34900 or other systems, applications, etc. with which the platform 34900 interacts. In aspects, the opportunity miners 35046 may be configured or used by developers of AI or RPA solutions to find opportunities for better solutions and to optimize existing solutions. In certain implementations, the opportunity mining system 35046 may include a set of systems that collect information within the platform 34900 and collect information within, about, and for a set of industrial environments and entities 34930, where the collected information has the potential to help identify and prioritize opportunities for increased automation and/or intelligence. For example only, the opportunity mining system 35046 may include systems that observe clusters of workers by time, by type, and by location, such as using cameras, wearables, or other sensors, to identify labor-intensive areas and processes in a set of industrial environments. These may be presented, such as in a ranked or prioritized list, or in a visualization (such as a heat map showing dwell times of workers on a map of an environment or a heat map showing routes traveled by workers within an environment) to show places with high labor activity. In various implementations, the industrial analytics application 35028 may be used to identify which environments or activities would most benefit from automation for purposes of labor saving.

[1973] In additional or alternative implementations, the opportunity mining system 35046 can include systems to characterize the extent of domain-specific or entity-specific knowledge or expertise required to undertake an action, use a program, use a machine, or perform any task in a process, for example, by observing the identity, credentials, experience, and/or other characteristics of worker(s) involved in the given process. This may be of particular benefit in situations where very experienced workers are involved (such as in maintenance or re-build processes on large or complex machines, or fine-tuning of complex processes where accumulated experience is required for effective work), especially where the population of those workers may be scarce (such as due

to retirement or a dwindling supply of new workers having the same credentials). Thus, the opportunity mining system 35046 may collect and supply to an industrial analytics application 35028 (such as for prioritizing the development of automation such as RPA) data indicating what processes of or about an industrial entity 34930 are most intensively dependent on workers that have particular sets of experience or credentials (such as ones that have experience or credentials that are scarce or diminishing). The opportunity mining system 35046 may, for example, correlate aggregated data (including trend information) on worker ages, credentials, and/or experience (including by process type) with data on the processes in which those workers are involved (such as by tracking locations of workers by type, by tracking time spent on processes by worker type, or otherwise). A set of high value automation opportunities may be automatically recommended based on a ranking set, such as one that weights opportunities at least in part based on the relative dependence of a set of processes on workers who are scarce or are expected to become scarcer.

[1974] In various aspects, the opportunity mining system 35046 may use information relating to the cost of the workers involved in a set of processes, such as by accessing worker data 34922, including human resource database information indicating the salaries of various workers (either as individuals or by type), information about the rates charged by service workers or other contractors, or other form of cost data. The opportunity mining system 35046 may provide such cost information for correlation with process tracking information, such as to enable an industrial analytics application 35028 to identify what processes are occupying the most time of the most expensive workers. This may include visualization of such processes, such as by heat maps that show what locations, routes, or processes are involving the most expensive time of workers in industrial environments or with respect to industrial entities 34930. The opportunity mining system 35046 may supply a ranked list, weighted list, or other form of data set indicating to developers what areas are most likely to benefit from further automation or artificial intelligence deployment.

[1975] In certain aspects, the opportunity mining system 35046 may "mine" an industrial environment for RPA opportunities by searching a human resources database and/or other labor-tracking database for areas that involve labor-intensive processes. For example only, the opportunity mining system 35046 may search a system for areas where credentials of workers indicate a relatively high potential for automation, may track clusters of workers (e.g., via a wearable device or other sensor) to find labor-intensive machines or processes, and/or track clusters of workers (e.g., via a wearable device or other sensor) by type of worker to find labor-intensive processes.

[1976] The opportunity mining system 35046 may include facilities for solicitation of appropriate training data sets that may be used to facilitate process automation. Certain kinds of data or other inputs, if available, may provide very high value for automation, such as video data sets that capture very experienced and/or highly expert workers performing complex tasks. Thus, the opportunity mining system 35046 can search for such video data sets as described herein. In the absence of a successful search for such data, or to supplement available data, the platform 34900 may include systems by which a user, such as a developer, may specify a desired type of data, such as software interaction data (for example, of an expert working with a program to perform a particular task),

video data (such as video showing a set of experts performing a certain kind of repair, an expert rebuilding a machine, an expert optimizing a certain kind of complex process, or similar), and/or physical process observation data (such as video or other type of sensor data).

[1977] The platform 34900 may be used to solicit such data, such as by offering some form of consideration (a monetary reward, tokens, cryptocurrency, licenses or rights, revenue sharing, or other consideration) to parties that provide data of the requested type. Rewards may be provided to parties for supplying pre-existing data and/or for undertaking steps to capture expert interactions, such as by taking video of a process. The resulting library of interactions captured in response to specification, solicitation, and rewards may be captured as a data set in the data storage layer 34910, such as for consumption by various applications 34912, elements of the adaptive intelligence systems layer 34904, and other processes and systems. In aspects, the library may include videos that are specifically developed as instructional videos to, among other uses, facilitate developing an automation map that can follow instructions in the video, such as by providing a sequence of steps according to a procedure or protocol, by breaking down the procedure or protocol into sub-steps that are candidates for automation, and the like. For example only, such instructional videos may be processed by natural language processing, such as to automatically develop a sequence of labeled instructions that can be used by a developer to facilitate a map, a graph, or other model of a process that assists with the development of automation for the process. In aspects, a specified set of training data sets may be configured to operate as inputs to learning. For example only, the training data may be time-synchronized with other data within the platform 34900 (such as outputs and outcomes from applications 34912, outputs and outcomes of industrial entities 34930, or the like) so that a given video of a process can be associated with those outputs and outcomes, thereby enabling feedback on learning that is sensitive to the outcomes that occurred for a captured process.

[1978] Referring to Figure 206, a set of opportunity miners 35046 may be provided as part of the adaptive intelligence layer 34904, which may be configured to seek and recommend opportunities to improve one or more of the elements of the platform 34900, such as via addition of artificial intelligence 35048, automation (including robotic process automation 35046), or the like to one or more of the systems, sub-systems, components, applications or the like of the platform 100 or with which the platform 100 interacts. In example embodiments, the opportunity miners 35046 may be configured or used by developers of AI or RPA solutions to find opportunities for better solutions and to optimize existing solutions. In example embodiments, the opportunity miners 35046 may include a s set of systems that collect information within the platform 100 and collect information within, about and for a set of industrial environments and entities 34930, where the collected information has the potential to help identify and prioritize opportunities for increased automation and/or intelligence. For example, the opportunity miners 35046 may include systems that observe clusters of workers by time, by type, and by location, such as using cameras, wearables, or other sensors, such as to identify labor-intensive areas and processes in set of industrial environments. These may be presented, such as in a ranked or prioritized list, or in a visualization (such as a heat map showing dwell times of workers on a map of an environment or a heat map showing routes

traveled by workers within an environment) to show places with high labor activity. In example embodiments, analytics 35028 may be used to identify which environments or activities would most benefit from automation for purposes of labor saving.

[1979] In example embodiments, opportunity miners 35046 may include systems to characterize the extent of domain-specific or entity-specific knowledge or expertise required to undertake an action, use a program, use a machine, or the like, such as observing the identity, credentials and experience of workers involved in given processes. This may be of particular benefit in situations where very experienced workers are involved (such as in maintenance or re-build processes on large or complex machines, or fine-tuning of complex processes where accumulated experience is required for effective work), especially where the population of those workers may be scarce (such as due to retirement or a dwindling supply of new workers having the same credentials. Thus, a set of opportunity miners 35046 may collect and supply to an analytics solution 35028, such as for prioritizing the development of automation 35042, data indicating what processes of or about an industrial entity 34930 are most intensively dependent on workers that have particular sets of experience or credentials, such as ones that have experience or credentials that are scarce or diminishing. The opportunity miners 35046 may, for example, correlate aggregated data (including trend information) on worker ages, credentials, experience (including by process type) with data on the processes in which those workers are involved (such as by tracking locations of workers by type, by tracking time spent on processes by worker type, and the like). A set of high value automation opportunities may be automatically recommended based on a ranking set, such as one that weights opportunities at least in part based on the relative dependence of a set of processes on workers who are scarce or are expected to become more scarce.

[1980] In example embodiments, the set of opportunity miners 35046 may use information relating to the cost of the workers involved in a set of processes, such as by accessing worker data 34922, including human resource database information indicating the salaries of various workers (either as individuals or by type), information about the rates charged by service workers or other contractors, or the like. An opportunity miner 35046 may provide such cost information for correlation with process tracking information, such as to enable an analytics solution 35028 to identify what processes are occupying the most time of the most expensive workers. This may include visualization of such processes, such as by heat maps that show what locations, routes, or processes are involving the most expensive time of workers in industrial environments or with respect to industrial entities 34930. The opportunity miners 35046 may supply a ranked list, weighted list, or other data set indicating to developers what areas are most likely to benefit from further automation or artificial intelligence deployment.

[1981] In example embodiments, mining an industrial environment for robotic process automation opportunities may include searching an HR database and/or other labor-tracking database for areas that involve labor-intensive processes; searching a system for areas where credentials of workers indicating potential for automation; tracking clusters of workers by a wearable to find labor-intensive machines or processes; tracking clusters of workers by a wearable by type of worker to find labor-intensive processes, and the like.

[1982] In example embodiments, opportunity mining may include facilities for solicitation of appropriate training data sets that may be used to facilitate process automation. For example, certain kinds of inputs, if available, would provide very high value for automation, such as video data sets that capture very experienced and/or highly expert workers performing complex tasks. Opportunity miners 35046 may search for such video data sets as described herein; however, in the absence of success (or to supplement available data), the platform may include systems by which a user, such as a developer, may specify a desired type of data, such as software interaction data (such as of an expert working with a program to perform a particular task), video data (such as video showing a set of experts performing a certain kind of repair, an expert rebuilding a machine, an expert optimizing a certain kind of complex process, or the like), physical process observation data (such as video, sensor data, or the like). The specification may be used to solicit such data, such as by offering some form of consideration (e.g., monetary reward, tokens, cryptocurrency, licenses or rights, revenue share, or other consideration) to parties that provide data of the requested type. Rewards may be provided to parties for supplying pre-existing data and/or for undertaking steps to capture expert interactions, such as by taking video of a process. The resulting library of interactions captured in response to specification, solicitation and rewards may be captured as a data set in the data storage layer 34910, such as for consumption by various applications 34912, adaptive intelligence systems 34904, and other processes and systems. In example embodiments, the library may include videos that are specifically developed as instructional videos, such as to facilitate developing an automation map that can follow instructions in the video, such as providing a sequence of steps according to a procedure or protocol, breaking down the procedure or protocol into sub-steps that are candidates for automation, and the like. In example embodiments, such videos may be processed by natural language processing, such as to automatically develop a sequence of labeled instructions that can be used by a developer to facilitate a map, a graph, or other model of a process that assists with development of automation for the process. In embodiments a specified set of training data sets may be configured to operate as inputs to learning. In such cases the training data may be time-synchronized with other data within the platform 34900, such as outputs and outcomes from applications 34912, outputs and outcomes of industrial entities 34930, or the like, so that a given video of a process can be associated with those outputs and outcomes, thereby enabling feedback on learning that is sensitive to the outcomes that occurred when a given process that was captured (such as on video, or through observation of software interactions or physical process interactions).

[1983] As noted elsewhere herein and in documents incorporated by reference, artificial intelligence (such as any of the techniques or systems described throughout this disclosure) may, in connection with various industrial entities 34930, functions and applications, be used to facilitate, among other things: (a) the optimization, automation and/or control of various functions, workflows, applications, features, resource utilization and other factors, (b) recognition or diagnosis of various states, entities, patterns, events, contexts, behaviors, or other elements; and/or (c) the forecasting of various states, events, contexts or other factors. As artificial intelligence improves, a large array of domain-specific and/or general artificial intelligence systems have

become available and are likely to continue to proliferate. As developers seek solutions to domain-specific problems, such as ones relevant to industrial entities 34930 and various applications of the platform 34902 described throughout this disclosure they face challenges in selecting artificial intelligence models (such as what set of neural networks, machine learning systems, expert systems, or the like to select) and in discovering and selecting what inputs may enable effective and efficient use of artificial intelligence for a given problem. As noted above, opportunity miners 35046 may assist with the discovery of opportunities for increased automation and intelligence; however, once opportunities are discovered, selection and configuration of an artificial intelligence solution still presents a significant challenge, one that is likely to continue to grow as artificial intelligence solutions proliferate.

[1984] One set of solutions to these challenges is an artificial intelligence store 34904 that is configured to enable collection, organization, recommendation and presentation of relevant sets of artificial intelligence systems based on one or more attributes of a domain and/or a domain-related problem. In example embodiments, an artificial intelligence store 34904 may include a set of interfaces to artificial intelligence systems, such as enabling the download of relevant artificial intelligence applications, establishment of links or other connections to artificial intelligence systems (such as links to cloud-deployed artificial intelligence systems via APIs, ports, connectors, or other interfaces) and the like. The artificial intelligence store 34904 may include descriptive content with respect to each of a variety of artificial intelligence systems, such as metadata or other descriptive material indicating suitability of a system for solving particular types of problems (e.g., forecasting, NLP, image recognition, pattern recognition, motion detection, route optimization, or many others) and/or for operating on domain-specific inputs, data or other entities. In example embodiments, the artificial intelligence store 34904 may be organized by category, such as domain, input types, processing types, output types, computational requirements and capabilities, cost, energy usage, and other factors. In example embodiments, an interface to the application store 34904 may take input from a developer and/or from the platform (such as from an opportunity miner 35046) that indicates one or more attributes of a problem that may be addressed through artificial intelligence and may provide a set of recommendations, such as via an artificial intelligence attribute search engine, for a subset of artificial intelligence solutions that may represent favorable candidates based on the developer's domain-specific problem. Search results or recommendations may, in example embodiments, be based at least in part on collaborative filtering, such as by asking developers to indicate or select elements of favorable models, as well as by clustering, such as by using similarity matrices, k-means clustering, or other clustering techniques that associate similar developers, similar domain-specific problems, and/or similar artificial intelligence solutions. The artificial intelligence store 34904 may include e-commerce features, such as ratings, reviews, links to relevant content, and mechanisms for provisioning, licensing, delivery and payment (including allocation of payments to affiliates and or contributors), including ones that operate using smart contract and/or blockchain features to automate purchasing, licensing, payment tracking, settlement of transactions, or other features.

[1985] In example embodiments, another set of solutions, which may be deployed alone or in connection with other elements of the platform, including the artificial intelligence store 34904, may include a set of functional imaging capabilities 34902, which may comprise monitoring systems 34906 and in some cases physical process observation systems 35058 and/or software interaction observation systems 35050, such as for monitoring various industrial entities 34930. Functional imaging systems 34902 may, in example embodiments, provide considerable insight into the types of artificial intelligence that are likely to be most effective in solving particular types of problems most effectively. As noted elsewhere in this disclosure and in the documents incorporated by reference herein, computational and networking systems, as they grow in scale, complexity and interconnections, manifest problems of information overload, noise, network congestion, energy waste, and many others. As the Internet of Things grows to hundreds of billions of devices, and virtually countless potential interconnections, optimization becomes exceedingly difficult. One source for insight is the human brain, which faces similar challenges and has evolved, over millennia, reasonable solutions to a wide range of very difficult optimization problems. The human brain operates with a massive neural network organized into interconnected modular systems, each of which has a degree of adaptation to solve particular problems, from regulation of biological systems and maintenance of homeostasis, to detection of a wide range of static and dynamic patterns, to recognition of threats and opportunities, among many others. Functional imaging 34902, such as functional magnetic resonance imaging (fMRI), electroencephalogram (EEG), computed tomography (CT) and other brain imaging systems have improved to the point that patterns of brain activity can be recognized in real time and temporally associated with other information, such behaviors, stimulus information, environmental condition data, gestures, eye movements, and other information, such that via functional imaging 34902, either alone or in combination with other information collected by monitoring systems 34906, the platform may determine and classify what brain modules, operations, systems, and/or functions are employed during the undertaking of a set of tasks or activities, such as ones involving software interaction 35050, physical process observations 35058, or a combination thereof. This classification may assist in selection and/or configuration of a set of artificial intelligence solutions, such as from an artificial intelligence store 34904, that includes a similar set of capabilities and/or functions to the set of modules and functions of the human brain when undertaking an activity, such as for the initial configuration of a robotic process automation (RPA) system 35042 that automates a task performed by an expert human. Thus, the platform may include a system that takes input from a functional imaging system 34902 to configure, optionally automatically based on matching of attributes between one or more biological systems, such as brain systems, and one or more artificial intelligence systems, a set of artificial intelligence capabilities for a robotic process automation system. Selection and configuration may further comprise selection of inputs to robotic process automation and/or artificial intelligence that are configured at least in part based on functional imaging of the brain while workers undertake tasks, such as selection of visual inputs (such as images from cameras) where vision systems of the brain are highly activated, selection of acoustic inputs where auditory systems of the brain are highly activated, selection of chemical inputs (such

as chemical sensors) where olfactory systems of the brain are highly activated, or the like. Thus, a biologically aware robotic process automation system may be improved by having initial configuration, or iterative improvement, be guided, either automatically or under developer control, by imaging-derived information collected as workers perform expert tasks that may benefit from automation.

[1986] Referring to Figure 207, additional details of an embodiment of the platform 34900 are provided, in particular relating to elements of the adaptive intelligent systems layer 34904 that facilitate improved edge intelligence, including the adaptive edge compute management system 35030 and the edge intelligence system 35038. These elements provide a set of systems that adaptively manage "edge" computation, storage and processing, such as by varying storage locations for data and processing locations (e.g., optimized by AI) between on-device storage, local systems, in the network, and in the cloud. The adaptive edge compute management system 35030 and the edge intelligence system 35038 enable facilitation of a dynamic definition by a user, such as a developer, operator, or host of the platform 100, of what constitutes the "edge" for purposes of a given application. For example only, for environments where data connections are slow or unreliable such as where an industrial facility does not have good access to cellular networks (e.g., due to remoteness of some environments (such as for drilling, construction, pipelining, or exploration), shielding or interference (such as where thick concrete or presence of large metal equipment interferes with networking performance), and/or congestion (such as where there are many devices seeking access to limited networking facilities)), edge computing capabilities can be defined and deployed to operate on the local area network of an environment, in peer-to-peer networks of devices, or on computing capabilities of local industrial entities 34930. Where strong data connections are available (such as where good backhaul facilities exist), edge computing capabilities can be disposed in the network, such as for caching frequently used data at locations that improve input/output performance, reduce latency, or otherwise improve performance of the platform 34900. Thus, adaptive definition and specification of where edge computing operations is enabled. This adaptive definition/specification can be under control of a developer or operator and/or determined automatically (such as by an expert system or automation system, e.g., based on detected network conditions for an environment, for an industrial entity 34930, or for a network as a whole). In certain implementations, the edge intelligence system 35038 can enable adaptation of edge computation (where computation occurs within various available networking resources, how networking occurs (e.g., by protocol selection), where data storage occurs, etc.) that is multi-application aware, such as accounting for QoS, latency requirements, congestion, and cost as understood and prioritized based on awareness of the requirements, the prioritization, and the value (including ROI, yield, and cost information, such as costs of failure) of edge computation capabilities across more than one application, including any combinations and subsets of the applications 34912 described herein or in the documents incorporated herein by reference.

[1987] In various aspects, the edge intelligence system 35038 can be enabled in part by edge computation capabilities, such as using a tensor processing unit (TPU), such as a single-board computing device running an edge-based Tensor Processing Unit (TPU) from Google™. In

additional or alternative aspects, the edge intelligence system 35038 can use a system-on-module (SOM) capability, such as a Coral™ SOM, as well as one or more accessories that are configured to provide machine learning inferencing capabilities to edge devices and systems, e.g., USB-connected accessories, Power-over-Ethernet (PoE) powered accessories, and accessories connected via other local power and data protocols. Such capabilities for edge intelligence system 35038 can be deployed in edge devices and systems of or about various industrial entities 34930 and may be used to provide pattern recognition, prediction, inferencing, and the like for various purposes, such as for predictive maintenance, recommendation of service and repairs, anomaly detection, fault detection, recognition of process failures, process optimization, machine vision, visual inspection, robotics, process automation, status reporting, natural language processing, diagnostic condition recognition, and voice recognition.

[1988] For example only, the edge TPU may include an application-specific integrated circuit (ASIC) and may feature, for example, an NXP™ i.MX 8M system-on-chip (SOC), a quad-core Cortex-A53 and a Cortex-M4F, or similar processing device. The system can, for example, use a graphics GPU, such as an integrated GC7000 Lite Graphics GPU, with RAM (e.g., 1GB of RAM) and Flash memory (e.g., 8GB or more of Flash memory).

[1989] In implementations, the system may include a variety of ports to enable linking of edge intelligence capability to various edge devices and systems via various protocols, such as via a MicroSD slot, a Gigabit or other Ethernet port, PoE ports, and various audio ports. Various wireless protocols may be supported, including NFC, WiFi, ZIGBEE® and BLUETOOTH™ 4.1. Connectivity may include wired connectivity such as USB connectivity, such as via Type-C OTG, a Type-C power connection, a Type-A 3.0 host, and/or a micro-B serial console. In aspects, the SOM can be integrated into an edge device or system, such as a Raspberry Pi or other Linux system, or a system using another conventional operating system. In further aspects, elements of the system can run a software operating system, such as a Linux-based system, such as Mendel™. Further, in certain implementations, models using an AI modeling system, such as TensorFlow™, can be compiled to run on the system.

[1990] Referring to Figure 208, additional details, components, sub-systems, and other elements of an example implementation of the industrial entity-oriented data storage systems layer 34910 of the platform 34900 are illustrated, relating in particular to implementations that include a geofenced virtual asset tag 35088. The virtual asset tag 35088 can be implemented as a data structure that contains data about an industrial entity 34930 (a machine, item of equipment, item of inventory, manufactured article, component, tool, device, worker, etc.), where the data is intended to be "tagged" to the asset. For example only, the data can relate uniquely to the particular asset (e.g., to a unique identifier for the individual asset) and can be linked to proximity to the asset (such as being geofenced to an area or location of or near the asset). The virtual asset tag 35088 is thus functionally equivalent to a physical asset tag, such as an RFID tag, in that it provides a local reader or similar device access to the data structure (as a reader would access an RFID tag) when the local reader or similar device is in proximity to the virtual asset stage 35088. In some aspects, access control can be managed and/or controlled as if the tag were physically located on an asset.

For example only, certain data may be encrypted with keys that only permit it to be read, written to, modified, etc. by an operator who is verified to be in the proximity of a tagged industrial entity 34930. In this implementation, partitioning of local-only data processing from remote data processing can be enabled.

[1991] In some aspects, the virtual asset tag 35088 can be configured to recognize the presence of an RF reader or other reader (such as by recognition of an interrogation signal) and communicate with the reader (such as with the help of protocol adaptors), e.g., over an RF communication link or other communication protocol, notwithstanding the absence of a conventional RFID tag. This may occur by communications from IoT devices, telematics systems, and by other devices residing on a local area network. In additional or alternative embodiments, a set of IoT devices in an industrial environment can act as distributed blockchain nodes, such as for storage of virtual asset tag data, for tracking of transactions, and for validation (such as by various consensus protocols) of enchained data, including transaction history for maintenance, repair, and service. The IoT devices in a geofence can collectively validate location and identity of a fixed asset that is tagged by a virtual asset tag 35088, such as where peers or neighbors validate other peers or neighbors as being in a given location, thereby validating the unique identity and location of the asset. Validation can use voting protocols, consensus protocols, other protocols, or combinations thereof. In aspects, the identity of the industrial entities 34930 that are tagged can be maintained in a blockchain. Additionally or alternatively, in some aspects, a virtual asset tag 35088 can include information that is related to an industrial digital thread 35084, such as historical information about an asset, its components, its history, etc.

[1992] Referring to Figure 209, in various aspects, the RPA system 35042 can be configured for developing and deploying one or more automation capabilities, including or enabling capabilities for a robot operational analytics system 35502. The robot operational analytics system 35502 can, in certain aspects, analyze operational actions of a set of robots, including with respect to location, mobility, and routing of mobile robots, as well as with respect to motions of robot components, such as where robots and/or robotic components are used within a wide range of protocols or procedures (such as manufacturing processes, assembly processes, transport processes, maintenance and repair processes, data collection processes).

[1993] In aspects, the RPA system 35042 may include or enable capabilities for machine learning on unstructured data 35508, including but not limited learning on a training set of human labels, tags, or other activities that allow characterization of the unstructured data, extraction of content from unstructured data, and/or generation of diagnostic codes or similar summaries from content of unstructured data. For example only, the RPA system 35042 may include sub-systems or capabilities for processing technical documents (such as technical data sheets, functional specifications, repair instructions, user manuals, and other documentation about industrial entities 34930), for processing human-entered notes (such as notes involved in diagnosis of problems, notes involved in prescribing or recommending actions, notes involved in characterizing operational activities, and notes involved in maintenance and repair operations), for processing information such as unstructured content contained on websites, social media feeds, etc. (such as

information about products or systems in an industrial environment that can be obtained from vendor websites), and other documentation.

[1994] In certain aspects, the RPA system 35042 may comprise a unified platform with a set of RPA capabilities, as well as system(s) for monitoring (such as the systems of the monitoring layer 34906 and data collection systems 34918), raw data processing system(s) 35504 (including but not limited to systems for optical character recognition (OCR), natural language processing (NPL), computer vision processing, sound processing, and other forms of sensor processing); workflow characterization and management system(s) 35516; analytics system(s) 35510; artificial intelligence system(s) 35048; and administrative system(s) 35514 (such as for policy, governance, and provisioning of services, roles, access controls, etc. In certain implementations, the RPA system 35042 can include such capabilities as a set of microservices in a microservices architecture. The RPA system 35042 may have a set of interfaces to other platform layers 34908, as well as to external systems, for data exchange such that the RPA system 35042 can be accessed as an RPA platform-as-a-service by other platform layers 34908 and/or external systems that can benefit from one or more automation capabilities.

[1995] In example embodiments, the RPA system 35042 may include a quality-of-work characterization system 35512 that can, e.g., identify high quality work as compared to other work or otherwise rate, gauge, or characterize work quality. Examples of such characterization of work quality services include recognizing human work as different from work performed by machines, recognizing which human work is likely to be of highest quality (such as work involving the most experienced or expensive personnel), recognizing which machine-performed work is likely to be of the highest quality (such as work that is performed by machines that have extensively learned on feedback from many outcomes, as compared to machines that are newly deployed), and recognizing which work has historically provided favorable outcomes (such as based on analytics or correlation to past outcomes). A set of thresholds may be applied, which may be varied under control of a developer or other user of the RPA system 35042, to indicate by type, by quality-level, or other measurement, which data sets indicating past work will be used for training within the machine learning systems that facilitate automation in the RPA system 35042.

[1996] As briefly mentioned above, a set of protocol adaptors can facilitate adaptive protocol transformations of data within the IIoT system. With reference to Figures 210-212, an example method and system for data processing in an industrial environment that utilizes protocol adaptors is illustrated in conjunction with the various components, interfaces, machines, devices, programs, methods, processes, protocols, and other elements collectively referred to herein as a platform 35600. In various implementations, the platform 35600 may include an intelligent, automated, machine learning, or otherwise "smart" protocol adaptor (referred to herein except where context indicates otherwise as a self-organizing protocol adaptor 35602) that may connect to one or more cloud, networked, and/or distributed computing platforms (referred to herein except where context indicates otherwise as IoT cloud platforms 35610).

[1997] The platform 35600 may include, connect to, or integrate with one or more sensors 35622 that may connect to the self-organizing protocol adaptor 35602 or to one or more IoT cloud

platforms 35610. In this manner, the one or more sensors 35622 can provide information about the industrial environment, about one or more machines, components, or devices in the industrial environment, about one or more network conditions (such as network bandwidth, spectrum availability, congestion, interference, cost, timing, and/or availability), or about one or more cloud conditions or parameters. Among other things, the sensors 35622 may be used by the self-organizing protocol adaptor 35602 to facilitate organization or selection of an appropriate protocol by which one or more IoT devices (such as an industrial IoT device 35620 in an industrial environment 35624) can communicate. The platform 35600 may include one or more external data sources 35618 (such as databases, data warehouses, data streams, data packages, mobile data collectors, or other sources) that are located in the industrial environment 35624 or elsewhere, including in the cloud 35612. Various IoT devices 35620 can be located in the industrial environment 35624. In some aspects, an IoT cloud platform 35610 is deployed in the cloud 35612 and has one or more interfaces 35614 by which various networked devices, such as the industrial IoT devices 35620, can connect to the IoT cloud platform 35610 via one or more protocols 35608.

[1998] In aspects, the sensors 35612 may include one or more of touch ID, chemical, electrical, acoustic, vibration, acceleration, velocity, position, light, motion, temperature, magnetic fields, gravity, humidity, moisture, pressure, electrical fields, and sound sensors.

[1999] The self-organizing protocol adaptor 35602 can select, create, determine, and/or organize a self-organizing protocol, which can be at least one of a centralized protocol, a distributed protocol, and a hybrid protocol. In some aspects, the self-organizing protocol is self-organized by artificial intelligence, e.g., via at least one of an expert system, a machine learning system, a deep learning system, and a neural network to select, create, determine, and/or organize the self-organizing protocol. For example only, the IoT cloud platform 35610 can use one or more protocols 35608 selected from the group consisting of REST/HTTP, websockets, MQTT, CoAP, M2M IoT, Modbus, XMPP, and DDS, although any protocol that is suitable for use is within the scope of the present disclosure.

[2000] In some implementations, the IoT cloud platform 35610 is at least one of a Websphere platform, an AWS platform, an Azure platform, a Google cloud platform, an IBM Watson platform, an Oracle platform, an SAP platform, a GE Predix platform, a Cisco platform, and a Bosch platform. It should be appreciated, however, that the IoT cloud platform 35610 can be of any type or form. Further, in various aspects, the industrial IoT device 35620 may be one or more of internet protocol (IP) capable devices, non-IP capable devices, IoT client devices, low power devices, java devices, or any other suitable IoT device.

[2001] In various aspects, the industrial environment 35624 is one or more of an energy production environment, a manufacturing environment, an energy extraction environment, and a construction environment.

[2002] In additional or alternative implementations, methods and systems are provided for industrial data processing having a self-organizing protocol adaptor 35602 and having a smart industrial heater 35604.

[2003] In additional or alternative implementations, an IoT cloud platform 35610 may include an IoT data adaptor 35700. The IoT data adaptor 35700, as depicted in Figure 211, may receive IoT data 35710 as an input. Input can be received from any one or more than one of the many external data sources 35618 identified elsewhere in this disclosure (such as databases, data warehouses, data streams, data packages, and mobile data collectors), sensors 35622, and any other data source. In some implementations, the IoT data adaptor 35700 can establish a connection to publish the data to one or more available IoT cloud platforms 35610, or to any other device, server computing device, etc. capable of receiving data. In some aspects, the connection can also or alternatively be established to one or more available IoT cloud platforms 35610 by detecting conditions, e.g., with a condition detector 35716. The conditions can be related to the connection attempt or attempts made by the IoT data adapter 35700 to one or more IoT cloud platforms 35610. These conditions related to the attempt or attempts can include the receipt of reply messages 35718 from an IoT cloud platform 35610. The reply messages 35718 can indicate connection success or failure and/or can include content that suggests alternative protocols that might result in a successful connection being established or similar content, such as data from the cloud platform or usage indicators.

[2004] In some aspects, the data received from the IoT adapter 35700 by the IoT cloud platform 35610 can be published by the IoT cloud platform 35610 by automatically formatting, wrapping, translating, or otherwise preparing a data package 35720 or data stream 35722. The data package 35720 or data stream 35722 can be formatted in any one of the wide range of available data formats, such as, but not limited to, those described elsewhere in this disclosure.

[2005] Optionally, the IoT data adapter 35700 can include an adaptation engine 35724 for the implementation of the adaptation techniques described herein. The IoT data adapter 35700 can use adaptation techniques to establish a successful connection to one or more than one IoT cloud platforms 35610. The adaptation techniques can include using any of the machine-learning techniques described elsewhere in this disclosure.

[2006] The IoT data adaptor 35700, in various aspects, can also or alternatively make connections from a data marketplace. In such implementations, a data package 35720 related to a first connection of a new data source may prompt a user interface of an IoT cloud platform 35610 with a message that indicates the availability of a new data source, how to integrate the data source (for example by providing metadata about the data source and/or the terms for using the data), and other similar information.

[2007] With specific reference to Figure 212, an example connect attempt according to some aspects of the present disclosure is depicted. A sensor swarm 35810 attempts to establish an HTTP protocol connection 35814 to an IoT cloud platform 35610, through a condition detector 35716. The IoT cloud platform 35610 rejects the attempt to establish the HTTP protocol connection 35814 and sends a reply message 35718 to the IoT data adapter 35700 indicating the attempt failure. Upon receipt of the message indicating failure of the attempt to establish the HTTP protocol connection 35814, the adaptation engine 35724 can send a message to the sensor swarm 35810, through the condition detector 35716, indicative of the failure. Further, in some aspects, the message from the adaptation engine 35724 can include information relating to a suggestion that the sensor swarm

35810 retry the connection to the IoT cloud platform 35610 using a different protocol, such as the illustrated MQTT protocol connection 35812. It should be appreciated that the example connect attempt illustrated in Figure 212 is merely illustrative and other connect attempts can include additional or fewer, or different, elements, messages, data, etc.

[2008] FIG. 213 illustrates an example environment of a digital twin system 40000. In example embodiments, the digital twin system 40000 generates a set of digital twins of a set of industrial environments 40020 and/or industrial entities within the set of industrial environments. In example embodiments, the digital twin system 40000 maintains a set of states of the respective industrial environments 40020, such as using sensor data obtained from respective sensor systems 40030 that monitor the industrial environments 40020. In example embodiments, the digital twin system 40000 may include a digital twin management system 40002, a digital twin I/O system 40004, a digital twin simulation system 40006, a digital twin dynamic model system 40008, a cognitive intelligence system 40010, and/or an environment control module 40012. In example embodiments, the digital twin system 40000 may provide a real time sensor API that provides a set of capabilities for enabling a set of interfaces for the sensors of the respective sensor systems 40030. In example embodiments, the digital twin system 40000 may include and/or employ other suitable APIs, brokers, connectors, bridges, gateways, hubs, ports, routers, switches, data integration systems, peer-to-peer systems, and the like to facilitate the transferring of data to and from the digital twin system 40000. In these embodiments, these connective components may allow an IIOT sensor or an intermediary device (e.g., a relay, an edge device, a switch, or the like) within a sensor system 40030 to communicate data to the digital twin system 40030 and/or to receive data (e.g., configuration data, control data, or the like) from the digital twin system 40030 or another external system. In example embodiments, the digital twin system 40000 may further include a digital twin datastore 40016 that stores digital twins 40018 of various industrial environments 40020 and the objects 40022, devices 40024, sensors 40026, and/or humans 40028 in the environment 40020.

[2009] A digital twin may refer to a digital representation of one or more industrial entities, such as an industrial environment 40020, a physical object 40022, a device 40024, a sensor 40026, a human 40028, or any combination thereof. Examples of industrial environments 40020 include, but are not limited to, a factory, a power plant, a food production facility (which may include an inspection facility), a commercial kitchen, an indoor growing facility, a natural resources excavation site (e.g., a mine, an oil field, etc.), and the like. Depending on the type of environment, the types of objects, devices, and sensors that are found in the environments will differ. Non-limiting examples of physical objects 40022 include raw materials, manufactured products, excavated materials, containers (e.g., boxes, dumpsters, cooling towers, vats, pallets, barrels, palates, bins, and the like), furniture (e.g., tables, counters, workstations, shelving, etc.), and the like. Non-limiting examples of devices 40024 include robots, computers, vehicles (e.g., cars, trucks, tankers, trains, forklifts, cranes, etc.), machinery/equipment (e.g., tractors, tillers, drills, presses, assembly lines, conveyor belts, etc.), and the like. The sensors 40026 may be any sensor devices and/or sensor aggregation devices that are found in a sensor system 40030 within an

environment. Non-limiting examples of sensors 40026 that may be implemented in a sensor system 40030 may include temperature sensors 40032, humidity sensors 40034, vibration sensors 40036, LIDAR sensors 40038, motion sensors 40040, chemical sensors 40042, audio sensors 40044, pressure sensors 40046, weight sensors 40048, radiation sensors 40050, video sensors 40052, wearable devices 40054, relays 40056, edge devices 40058, crosspoint switches 40060, and/or any other suitable sensors. Examples of different types of physical objects 40022, devices 40024, sensors 40026, and environments 40020 are referenced throughout the disclosure.

[2010] In example embodiments, a crosspoint switch 40060 is implemented in the sensor system 40030 having multiple inputs and multiple outputs including a first input connected to the first sensor and a second input connected to the second sensor. The multiple outputs include a first output and second output configured to be switchable between a condition in which the first output is configured to switch between delivery of the first sensor signal and the second sensor signal and a condition in which there is simultaneous delivery of the first sensor signal from the first output and the second sensor signal from the second output. Each of multiple inputs is configured to be individually assigned to any of the multiple outputs. Unassigned outputs are configured to be switched off producing a high-impedance state.

[2011] In example embodiments, the first sensor signal and the second sensor signal are continuous vibration data about the industrial environment. In example embodiments, the second sensor in the sensor system 40030 is configured to be connected to the first machine. In example embodiments, the second sensor in the sensor system 40030 is configured to be connected to a second machine in the industrial environment. In example embodiments, the computing environment of the platform is configured to compare relative phases of the first and second sensor signals. In example embodiments, the first sensor is a single-axis sensor and the second sensor is a three-axis sensor. In example embodiments, at least one of the multiple inputs of the crosspoint switch 40060 includes internet protocol, front-end signal conditioning, for improved signal-to-noise ratio. In example embodiments, the crosspoint switch 40060 includes a third input that is configured with a continuously monitored alarm having a pre-determined trigger condition when the third input is unassigned to any of the multiple outputs.

[2012] In example embodiments, multiple inputs of the crosspoint switch 40060 includes a third input connected to the second sensor and a fourth input connected to the second sensor. The first sensor signal is from a single-axis sensor at an unchanging location associated with the first machine. In example embodiments, the second sensor is a three-axis sensor. In example embodiments, the sensor system 40030 is configured to record gap-free digital waveform data simultaneously from at least the first input, the second input, the third input, and the fourth input. In example embodiments, the platform is configured to determine a change in relative phase based on the simultaneously recorded gap-free digital waveform data. In example embodiments, the second sensor is configured to be movable to a plurality of positions associated with the first machine while obtaining the simultaneously recorded gap-free digital waveform data. In example embodiments, multiple outputs of the crosspoint switch include a third output and fourth output. The second, third, and fourth outputs are assigned together to a sequence of tri-axial sensors each

located at different positions associated with the machine. In example embodiments, the platform is configured to determine an operating deflection shape based on the change in relative phase and the simultaneously recorded gap-free digital waveform data.

[2013] In example embodiments, the unchanging location is a position associated with the rotating shaft of the first machine. In example embodiments, tri-axial sensors in the sequence of the tri-axial sensors are each located at different positions on the first machine but are each associated with different bearings in the machine. In example embodiments, tri-axial sensors in the sequence of the tri-axial sensors are each located at similar positions associated with similar bearings but are each associated with different machines. In example embodiments, the sensor system 40030 is configured to obtain the simultaneously recorded gap-free digital waveform data from the first machine while the first machine and a second machine are both in operation. In example embodiments, the sensor system 40030 is configured to characterize a contribution from the first machine and the second machine in the simultaneously recorded gap-free digital waveform data from the first machine. In example embodiments, the simultaneously recorded gap-free digital waveform data has a duration that is in excess of one minute.

[2014] In example embodiments, a method of monitoring a machine having at least one shaft supported by a set of bearings includes monitoring a first data channel assigned to a single-axis sensor at an unchanging location associated with the machine. The method includes monitoring second, third, and fourth data channels each assigned to an axis of a three-axis sensor. The method includes recording gap-free digital waveform data simultaneously from all of the data channels while the machine is in operation and determining a change in relative phase based on the digital waveform data.

[2015] In example embodiments, the tri-axial sensor is located at a plurality of positions associated with the machine while obtaining the digital waveform. In example embodiments, the second, third, and fourth channels are assigned together to a sequence of tri-axial sensors each located at different positions associated with the machine. In example embodiments, the data is received from all of the sensors simultaneously. In example embodiments, the method includes determining an operating deflection shape based on the change in relative phase information and the waveform data. In example embodiments, the unchanging location is a position associated with the shaft of the machine. In example embodiments, the tri-axial sensors in the sequence of the tri-axial sensors are each located at different positions and are each associated with different bearings in the machine. In example embodiments, the unchanging location is a position associated with the shaft of the machine. The tri-axial sensors in the sequence of the tri-axial sensors are each located at different positions and are each associated with different bearings that support the shaft in the machine.

[2016] In example embodiments, the method includes monitoring the first data channel assigned to the single-axis sensor at an unchanging location located on a second machine. The method includes monitoring the second, the third, and the fourth data channels, each assigned to the axis of a three-axis sensor that is located at the position associated with the second machine. The method also includes recording gap-free digital waveform data simultaneously from all of the data channels

from the second machine while both of the machines are in operation. In example embodiments, the method includes characterizing the contribution from each of the machines in the gap-free digital waveform data simultaneously from the second machine.

[2017] In example embodiments, on-device sensor fusion and data storage for industrial IoT devices is supported, including on-device sensor fusion and data storage for an industrial IoT device, where data from multiple sensors is multiplexed at the device for storage of a fused data stream. For example, pressure and temperature data may be multiplexed into a data stream that combines pressure and temperature in a time series, such as in a byte-like structure (where time, pressure, and temperature are bytes in a data structure, so that pressure and temperature remain linked in time, without requiring separate processing of the streams by outside systems), or by adding, dividing, multiplying, subtracting, or the like, such that the fused data can be stored on the device. Any of the sensor data types described throughout this disclosure, including vibration data, can be fused in this manner and stored in a local data pool, in storage, or on an IoT device, such as a data collector, a component of a machine, or the like.

[2018] In example embodiments, a set of digital twins may represent an entire organization, such as energy production organizations, oil and gas organizations, renewable energy production organizations, aerospace manufacturers, vehicle manufacturers, heavy equipment manufacturers, mining organizations, drilling organizations, offshore platform organizations, and the like. In these examples, the digital twins may include digital twins of one or more industrial facilities of the organization.

[2019] In example embodiments, the digital twin management system 40002 generates digital twins. A digital twin may be comprised of (e.g., via reference) other digital twins. In this way, a discrete digital twin may be comprised of a set of other discrete digital twins. For example, a digital twin of a machine may include digital twins of sensors on the machine, digital twins of components that make up the machine, digital twins of other devices that are incorporated in or integrated with the machine (such as systems that provide inputs to the machine or take outputs from it), and/or digital twins of products or other items that are made by the machine. Taking this example one step further, a digital twin of an industrial facility (e.g., a factory) may include a digital twin representing the layout of the industrial facility, including the arrangement of physical assets and systems in or around the facility, as well as digital assets of the assets within the facility (e.g., the digital twin of the machine), as well as digital twins of storage areas in the facility, digital twins of humans collecting vibration measurements from machines throughout the facility, and the like. In this second example, the digital twin of the industrial facility may reference the embedded digital twins, which may then reference other digital twins embedded within those digital twins.

[2020] In example embodiments, a digital twin may represent abstract entities, such as workflows and/or processes, including inputs, outputs, sequences of steps, decision points, processing loops, and the like that make up such workflows and processes. For example, a digital twin may be a digital representation of a manufacturing process, a logistics workflow, an agricultural process, a mineral extraction process, or the like. In these embodiments, the digital twin may include references to the industrial entities that are included in the workflow or process. The digital twin

of the manufacturing process may reflect the various stages of the process. In some of these embodiments, the digital twin system 40000 receives real-time data from the industrial facility (e.g., from a sensor system 40030 of the environment 40020) in which the manufacturing process takes place and reflects a current (or substantially current) state of the process in real-time.

[2021] In example embodiments, the digital representation may include a set of data structures (e.g., classes) that collectively define a set of properties of a represented physical object 40022, device 40024, sensor 40026, or environment 40020 and/or possible behaviors thereof. For example, the set of properties of a physical object 40022 may include a type of the physical object, the dimensions of the object, the mass of the object, the density of the object, the material(s) of the object, the physical properties of the material(s), the surface of the physical object, the status of the physical object, a location of the physical object, identifiers of other digital twins contained within the object, and/or other suitable properties. Examples of behavior of a physical object may include a state of the physical object (e.g., a solid, liquid, or gas), a melting point of the physical object, a density of the physical object when in a liquid state, a viscosity of the physical object when in a liquid state, a freezing point of the physical object, a density of the physical object when in a solid state, a hardness of the physical object when in a solid state, the malleability of the physical object, the buoyancy of the physical object, the conductivity of the physical object, a burning point of the physical object, the manner by which humidity affects the physical object, the manner by which water or other liquids affect the physical object, a terminal velocity of the physical object, and the like. In another example, the set of properties of a device may include a type of the device, the dimensions of the device, the mass of the device, the density of the density of the device, the material(s) of the device, the physical properties of the material(s), the surface of the device, the output of the device, the status of the device, a location of the device, a trajectory of the device, vibration characteristics of the device, identifiers of other digital twins that the device is connected to and/or contains, and the like. Examples of the behaviors of a device may include a maximum acceleration of a device, a maximum speed of a device, ranges of motion of a device, a heating profile of a device, a cooling profile of a device, processes that are performed by the device, operations that are performed by the device, and the like. Example properties of an environment may include the dimensions of the environment, the boundaries of the environment, the temperature of the environment, the humidity of the environment, the airflow of the environment, the physical objects in the environment, currents of the environment (if a body of water), and the like. Examples of behaviors of an environment may include scientific laws that govern the environment, processes that are performed in the environment, rules or regulations that must be adhered to in the environment, and the like.

[2022] In example embodiments, the properties of a digital twin may be adjusted. For example, the temperature of a digital twin, a humidity of a digital twin, the shape of a digital twin, the material of a digital twin, the dimensions of a digital twin, or any other suitable parameters may be adjusted. As the properties of the digital twin are adjusted, other properties may be affected as well. For example, if the temperature of an environment 40020 is increased, the pressure within the environment may increase as well, such as a pressure of a gas in accordance with the ideal gas law.

In another example, if a digital twin of a subzero environment is increased to above freezing temperatures, the properties of an embedded twin of water in a solid state (i.e., ice) may change into a liquid state over time.

[2023] Digital twins may be represented in a number of different forms. In example embodiments, a digital twin may be a visual digital twin that is rendered by a computing device, such that a human user can view digital representations of an environment 40020 and/or the physical objects 40022, devices 40024, and/or the sensors 40026 within an environment. In example embodiments, the digital twin may be rendered and output to a display device. In some of these embodiments, the digital twin may be rendered in a graphical user interface, such that a user may interact with the digital twin. For example, a user may "drill down" on a particular element (e.g., a physical object or device) to view additional information regarding the element (e.g., a state of a physical object or device, properties of the physical object or device, or the like). In example embodiments, the digital twin may be rendered and output in a virtual reality display. For example, a user may view a 3D rendering of an environment (e.g., using monitor or a virtual reality headset). While doing so, the user may view/inspect digital twins of physical assets or devices in the environment.

[2024] In example embodiments, a data structure of the visual digital twins (i.e., digital twins that are configured to be displayed in a 2D or 3D manner) may include surfaces (e.g., splines, meshes, polygons meshes, or the like). In example embodiments, the surfaces may include texture data, shading information, and/or reflection data. In this way, a surface may be displayed in a more realistic manner. In example embodiments, such surfaces may be rendered by a visualization engine (not shown) when the digital twin is within a field of view and/or when existing in a larger digital twin (e.g., a digital twin of an industrial environment). In these embodiments, the digital twin system 40000 may render the surfaces of digital objects, whereby a rendered digital twin may be depicted as a set of adjoined surfaces.

[2025] In example embodiments, a user may provide input that controls one or more properties of a digital twin via a graphical user interface. For example, a user may provide input that changes a property of a digital twin. In response, the digital twin system 40000 can calculate the effects of the changed property and may update the digital twin and any other digital twins affected by the change of the property.

[2026] In example embodiments, a user may view processes being performed with respect to one or more digital twins (e.g., manufacturing of a product, extracting minerals from a mine or well, a livestock inspection line, and the like). In these embodiments, a user may view the entire process or specific steps within a process.

[2027] In example embodiments, a digital twin (and any digital twins embedded therein) may be represented in a non-visual representation (or "data representation"). In these embodiments, a digital twin and any embedded digital twins exist in a binary representation but the relationships between the digital twins are maintained. For example, in example embodiments, each digital twin and/or the components thereof may be represented by a set of physical dimensions that define a shape of the digital twin (or component thereof). Furthermore, the data structure embodying the digital twin may include a location of the digital twin. In example embodiments, the location of

the digital twin may be provided in a set of coordinates. For example, a digital twin of an industrial environment may be defined with respect to a coordinate space (e.g., a Cartesian coordinate space, a polar coordinate space, or the like). In example embodiments, embedded digital twins may be represented as a set of one or more ordered triples (e.g., [x coordinate, y coordinate, z coordinates] or other vector-based representations). In some of these embodiments, each ordered triple may represent a location of a specific point (e.g., center point, top point, bottom point, or the like) on the industrial entity (e.g., object, device, or sensor) in relation to the environment in which the industrial entity resides. In example embodiments, a data structure of a digital twin may include a vector that indicates a motion of the digital twin with respect to the environment. For example, fluids (e.g., liquids or gases) or solids may be represented by a vector that indicates a velocity (e.g., direction and magnitude of speed) of the entity represented by the digital twin. In example embodiments, a vector within a twin may represent a microscopic subcomponent, such as a particle within a fluid, and a digital twin may represent physical properties, such as displacement, velocity, acceleration, momentum, kinetic energy, vibrational characteristics, thermal properties, electromagnetic properties, and the like.

[2028] In example embodiments, a set of two or more digital twins may be represented by a graph database that includes nodes and edges that connect the nodes. In some implementations, an edge may represent a spatial relationship (e.g., "abuts", "rests upon", "contains", and the like). In these embodiments, each node in the graph database represents a digital twin of an entity (e.g., an industrial entity) and may include the data structure defining the digital twin. In these embodiments, each edge in the graph database may represent a relationship between two entities represented by connected nodes. In some implementations, an edge may represent a spatial relationship (e.g., "abuts", "rests upon", "interlocks with", "bears", "contains", and the like). In example embodiments, various types of data may be stored in a node or an edge. In example embodiments, a node may store property data, state data, and/or metadata relating to a facility, system, subsystem, and/or component. Types of property data and state data will differ based on the entity represented by a node. For example, a node representing a robot may include property data that indicates a material of the robot, the dimensions of the robot (or components thereof), a mass of the robot, and the like. In this example, the state data of the robot may include a current pose of the robot, a location of the robot, and the like. In example embodiments, an edge may store relationship data and metadata data relating to a relationship between two nodes. Examples of relationship data may include the nature of the relationship, whether the relationship is permanent (e.g., a fixed component would have a permanent relationship with the structure to which it is attached or resting on), and the like. In example embodiments, an edge may include metadata concerning the relationship between two entities. For example, if a product was produced on an assembly line, one relationship that may be documented between a digital twin of the product and the assembly line may be "created by". In these embodiments, an example edge representing the "created by" relationship may include a timestamp indicating a date and time that the product was created. In another example, a sensor may take measurements relating to a state of a device, whereby one relationship between the sensor and the device may include "measured" and may

define a measurement type that is measured by the sensor. In this example, the metadata stored in an edge may include a list of N measurements taken and a timestamp of each respective measurement. In this way, temporal data relating to the nature of the relationship between two entities may be maintained, thereby allowing for an analytics engine, machine-learning engine, and/or visualization engine to leverage such temporal relationship data, such as by aligning disparate data sets with a series of points in time, such as to facilitate cause-and-effect analysis used for prediction systems.

[2029] In example embodiments, a graph database may be implemented in a hierarchical manner, such that the graph database relates a set of facilities, systems, and components. For example, a digital twin of a manufacturing environment may include a node representing the manufacturing environment. The graph database may further include nodes representing various systems within the manufacturing environment, such as nodes representing an HVAC system, a lighting system, a manufacturing system, and the like, all of which may connect to the node representing the manufacturing system. In this example, each of the systems may further connect to various subsystems and/or components of the system. For example, within the HVAC system, the HVAC system may connect to a subsystem node representing a cooling system of the facility, a second subsystem node representing a heating system of the facility, a third subsystem node representing the fan system of the facility, and one or more nodes representing a thermostat of the facility (or multiple thermostats). Carrying this example further, the subsystem nodes and/or component nodes may connect to lower level nodes, which may include subsystem nodes and/or component nodes. For example, the subsystem node representing the cooling subsystem may be connected to a component node representing an air conditioner unit. Similarly, a component node representing a thermostat device may connect to one or more component nodes representing various sensors (e.g., temperature sensors, humidity sensors, and the like).

[2030] In embodiments where a graph database is implemented, a graph database may relate to a single environment or may represent a larger enterprise. In the latter scenario, a company may have various manufacturing and distribution facilities. In these embodiments, an enterprise node representing the enterprise may connect to environment nodes of each respective facility. In this way, the digital twin system 40000 may maintain digital twins for multiple industrial facilities of an enterprise.

[2031] In example embodiments, the digital twin system 40000 may use a graph database to generate a digital twin that may be rendered and displayed and/or may be represented in a data representation. In the former scenario, the digital twin system 40000 may receive a request to render a digital twin, whereby the request includes one or more parameters that are indicative of a view that will be depicted. For example, the one or more parameters may indicate an industrial environment to be depicted and the type of rendering (e.g., "real-world view" that depicts the environment as a human would see it, an "infrared view" that depicts objects as a function of their respective temperature, an "airflow view" that depicts the airflow in a digital twin, or the like). In response, the digital twin system 40000 may traverse a graph database and may determine a configuration of the environment to be depicted based on the nodes in the graph database that are

related (either directly or through a lower level node) to the environment node of the environment and the edges that define the relationships between the related nodes. Upon determining a configuration, the digital twin system 40000 may identify the surfaces that are to be depicted and may render those surfaces. The digital twin system 40000 may then render the requested digital twin by connecting the surfaces in accordance with the configuration. The rendered digital twin may then be output to a viewing device (e.g., VR headset, monitor, or the like). In some scenarios, the digital twin system 40000 may receive real-time sensor data from a sensor system 40030 of an environment 40020 and may update the visual digital twin based on the sensor data. For example, the digital twin system 40000 may receive sensor data (e.g., vibration data from a vibration sensor 40036) relating to a motor and its set of bearings. Based on the sensor data, the digital twin system 40000 may update the visual digital twin to indicate the approximate vibrational characteristics of the set of bearings within a digital twin of the motor.

[2032] In scenarios where the digital twin system 40000 is providing data representations of digital twins (e.g., for dynamic modeling, simulations, machine learning), the digital twin system 40000 may traverse a graph database and may determine a configuration of the environment to be depicted based on the nodes in the graph database that are related (either directly or through a lower level node) to the environment node of the environment and the edges that define the relationships between the related nodes. In some scenarios, the digital twin system 40000 may receive real-time sensor data from a sensor system 40030 of an environment 40020 and may apply one or more dynamic models to the digital twin based on the sensor data. In other scenarios, a data representation of a digital twin may be used to perform simulations, as is discussed in greater detail throughout the specification.

[2033] In example embodiments, the digital twin system 40000 may execute a digital ghost that is executed with respect to a digital twin of an industrial environment. In these embodiments, the digital ghost may monitor one or more sensors of a sensor system 40030 of an industrial environment to detect anomalies that may indicate a malicious virus or other security issues.

[2034] As discussed, the digital twin system 40000 may include a digital twin management system 40002, a digital twin I/O system 40004, a digital twin simulation system 40006, a digital twin dynamic model system 40008, a cognitive intelligence system 40010, and/or an environment control module 40012.

[2035] In example embodiments, the digital twin management system 40002 creates new digital twins, maintains/updates existing digital twins, and/or renders digital twins. The digital twin management system 40002 may receive user input, uploaded data, and/or sensor data to create and maintain existing digital twins. Upon creating a new digital twin, the digital twin management system 40002 may store the digital twin in the digital twin datastore 40016. Creating, updating, and rendering digital twins are discussed in greater detail throughout the disclosure.

[2036] In example embodiments, the digital twin I/O system 40004 receives input from various sources and outputs data to various recipients. In example embodiments, the digital twin I/O system receives sensor data from one or more sensor systems 40030. In these embodiments, each sensor system 40030 may include one or more IoT sensors that output respective sensor data. Each sensor

may be assigned an IP address or may have another suitable identifier. Each sensor may output sensor packets that include an identifier of the sensor and the sensor data. In example embodiments, the sensor packets may further include a timestamp indicating a time at which the sensor data was collected. In example embodiments, the digital twin I/O system 40004 may interface with a sensor system 40030 via the real-time sensor connectivity 40014 such as a webhook, API, or the like. In these embodiments, one or more devices (e.g., sensors, aggregators, edge devices) in the sensor system 40030 may transmit the sensor packets containing sensor data to the digital twin I/O system 40004 via the webhook, etc. The digital twin I/O system may determine the sensor system 40030 that transmitted the sensor packets and the contents thereof, and may provide the sensor data and any other relevant data (e.g., time stamp, environment identifier/sensor system identifier, and the like) to the digital twin management system 40002.

[2037] In example embodiments, the digital twin I/O system 40004 may receive imported data from one or more sources. For example, the digital twin system 40000 may provide a portal for users to create and manage their digital twins. In these embodiments, a user may upload one or more files (e.g., image files, LIDAR scans, blueprints, and the like) in connection with a new digital twin that is being created. In response, the digital twin I/O system 40004 may provide the imported data to the digital twin management system 40002. The digital twin I/O system 40004 may receive other suitable types of data without departing from the scope of the disclosure.

[2038] In example embodiments, the digital twin simulation system 40006 is configured to execute simulations using the digital twin. For example, the digital twin simulation system 40006 may iteratively adjust one or more parameters of a digital twin and/or one or more embedded digital twins. In embodiments the digital twin simulation system 40006, for each set of parameters, executes a simulation based on the set of parameters and may collect the simulation outcome data resulting from the simulation. Put another way, the digital twin simulation system 40006 may collect the properties of the digital twin and the digital twins within or containing the digital twin used during the simulation as well as any outcomes stemming from the simulation. For example, in running a simulation on a digital twin of an indoor agricultural facility, the digital twin simulation system 40006 can vary the temperature, humidity, airflow, carbon dioxide and/or other relevant parameters and can execute simulations that output outcomes resulting from different combinations of the parameters. In another example, the digital twin simulation system 40006 may simulate the operation of a specific machine within an industrial facility that produces an output given a set of inputs. In example embodiments, the inputs may be varied to determine an effect of the inputs on the machine and the output thereof. In another example, the digital twin simulation system 40006 may simulate the vibration of a machine and/or machine components. In this example, the digital twin of the machine may include a set of operating parameters, interfaces, and capabilities of the machine. In example embodiments, the operating parameters may be varied to evaluate the effectiveness of the machine. The digital twin simulation system 40006 is discussed in further detail throughout the disclosure.

[2039] In example embodiments, the digital twin dynamic model system 40008 is configured to model one or more behaviors with respect to a digital twin of an environment. In example

embodiments, the digital twin dynamic model system 40008 may receive a request to model a certain type of behavior regarding an environment or a process and may model that behavior using a dynamic model, the digital twin of the environment or process, and sensor data collected from one or more sensors that are monitoring the environment or process. For example, an operator of a machine having bearings may wish to model the vibration of the machine and bearings to determine whether the machine and/or bearings can withstand an increase in output. In this example, the digital twin dynamic model system 40008 may execute a dynamic model that is configured to determine whether an increase in output would result in adverse consequences (e.g., failures, downtime, or the like). The digital twin dynamic model system 40008 is discussed in further detail throughout the disclosure.

[2040] In example embodiments, the cognitive processes system 40010 performs machine learning and artificial intelligence related tasks on behalf of the digital twin system. In example embodiments, the cognitive processes system 40010 may train any suitable type of model, including but not limited to various types of neural networks, regression models, random forests, decision trees, Hidden Markov models, Bayesian models, and the like. In example embodiments, the cognitive processes system 40010 trains machine learned models using the output of simulations executed by the digital twin simulation system 40006. In some of these embodiments, the outcomes of the simulations may be used to supplement training data collected from real-world environments and/or processes. In example embodiments, the cognitive processes system 40010 leverages machine learned models to make predictions, identifications, classifications and provide decision support relating to the real-world environments and/or processes represented by respective digital twins.

[2041] For example, a machine-learned prediction model may be used to predict the cause of irregular vibrational patterns (e.g., a suboptimal, critical, or alarm vibration fault state) for a bearing of an engine in an industrial facility. In this example, the cognitive processes system 40010 may receive vibration sensor data from one or more vibration sensors disposed on or near the engine and may receive maintenance data from the industrial facility and may generate a feature vector based on the vibration sensor data and the maintenance data. The cognitive processes system 40010 may input the feature vector into a machine-learned model trained specifically for the engine (e.g., using a combination simulation data and real-world data of causes of irregular vibration patterns) to predict the cause of the irregular vibration patterns. In this example, the causes of the irregular vibrational patterns could be a loose bearing, a lack of bearing lubrication, a bearing that is out of alignment, a worn bearing, the phase of the bearing may be aligned with the phase of the engine, loose housing, loose bolt, and the like.

[2042] In another example, a machine-learned model may be used to provide decision support to bring a bearing of an engine in an industrial facility operating at a suboptimal vibration fault level state to a normal operation vibration fault level state. In this example, the cognitive processes system 40010 may receive vibration sensor data from one or more vibration sensors disposed on or near the engine and may receive maintenance data from the industrial facility and may generate a feature vector based on the vibration sensor data and the maintenance data. The cognitive

processes system 40010 may input the feature vector into a machine-learned model trained specifically for the engine (e.g., using a combination simulation data and real-world data of solutions to irregular vibration patterns) to provide decision support in achieving a normal operation fault level state of the bearing. In this example, the decision support could be a recommendation to tighten the bearing, lubricate the bearing, re-align the bearing, order a new bearing, order a new part, collect additional vibration measurements, change operating speed of the engine, tighten housings, tighten bolts, and the like.

[2043] In another example, a machine-learned model may be used to provide decision support relating to vibration measurement collection by a worker. In this example, the cognitive processes system 40010 may receive vibration measurement history data from the industrial facility and may generate a feature vector based on the vibration measurement history data. The cognitive processes system 40010 may input the feature vector into a machine-learned model trained specifically for the engine (e.g., using a combination simulation data and real-world vibration measurement history data) to provide decision support in selecting vibration measurement locations.

[2044] In yet another example, a machine-learned model may be used to identify vibration signatures associated with machine and/or machine component problems. In this example, the cognitive processes system 40010 may receive vibration measurement history data from the industrial facility and may generate a feature vector based on the vibration measurement history data. The cognitive processes system 40010 may input the feature vector into a machine-learned model trained specifically for the engine (e.g., using a combination simulation data and real-world vibration measurement history data) to identify vibration signatures associated with a machine and/or machine component. The foregoing examples are non-limiting examples and the cognitive processes system 40010 may be used for any other suitable AI/machine-learning related tasks that are performed with respect to industrial facilities.

[2045] In example embodiments, the environment control system 40012 controls one or more aspects of industrial facilities. In some of these embodiments, the environment control system 40012 may control one or more devices within an industrial environment. For example, the environment control system 40012 may control one or more machines within an environment, robots within an environment, an HVAC system of the environment, an alarm system of the environment, an assembly line in an environment, or the like. In example embodiments, the environment control system 40012 may leverage the digital twin simulation system 40006, the digital twin dynamic model system 40008, and/or the cognitive processes system 40010 to determine one or more control instructions. In example embodiments, the environment control system 40012 may implement a rules-based and/or a machine-learning approach to determine the control instructions. In response to determining a control instruction, the environment control system 40012 may output the control instruction to the intended device within a specific environment via the digital twin I/O system 40004.

[2046] FIG. 214 illustrates an example digital twin management system 40002 according to some embodiments of the present disclosure. In example embodiments, the digital twin management

system 40002 may include, but is not limited to, a digital twin creation module 40064, a digital twin update module 40066, and a digital twin visualization module 40068.

[2047] In example embodiments, the digital twin creation module 40064 may create a set of new digital twins of a set of environments using input from users, imported data (e.g., blueprints, specifications, and the like), image scans of the environment, 3D data from a LIDAR device and/or SLAM sensor, and other suitable data sources. For example, a user (e.g., a user affiliated with an organization/customer account) may, via a client application 40070, provide input to create a new digital twin of an environment. In doing so, the user may upload 2D or 3D image scans of the environment and/or a blueprint of the environment. The user may also upload 3D data, such as taken by a camera, a LIDAR device, an IR scanner, a set of SLAM sensors, a radar device, an EMF scanner, or the like. In response to the provided data, the digital twin creation module 40064 may create a 3D representation of the environment, which may include any objects that were captured in the image data/detected in the 3D data. In example embodiments, the cognitive processes system 40072 may analyze input data (e.g., blueprints, image scans, 3D data) to classify rooms, pathways, equipment, and the like to assist in the generation of the 3D representation. In example embodiments, the digital twin creation module 40064 may map the digital twin to a 3D coordinate space (e.g., a Cartesian space having x, y, and z axes).

[2048] In example embodiments, the digital twin creation module 40064 may output the 3D representation of the environment to a graphical user interface (GUI). In some of these embodiments, a user may identify certain areas and/or objects and may provide input relating to the identified areas and/or objects. For example, a user may label specific rooms, equipment, machines, and the like. Additionally or alternatively, the user may provide data relating to the identified objects and/or areas. For example, in identifying a piece of equipment, the user may provide a make/model number of the equipment. In example embodiments, the digital twin creation module 40064 may obtain information from a manufacturer of a device, a piece of equipment, or machinery. This information may include one or more properties and/or behaviors of the device, equipment, or machinery. In example embodiments, the user may, via the GUI, identify locations of sensors throughout the environment. For each sensor, the user may provide a type of sensor and related data (e.g., make, model, IP address, and the like). The digital twin creation module 40064 may record the locations (e.g., the x, y, z coordinates of the sensors) in the digital twin of the environment. In example embodiments, the digital twin system 40000 may employ one or more systems that automate the population of digital twins. For example, the digital twin system 40000 may employ a machine vision-based classifier that classifies makes and models of devices, equipment, or sensors. Additionally or alternatively, the digital twin system 40000 may iteratively ping different types of known sensors to identify the presence of specific types of sensors that are in an environment. Each time a sensor responds to a ping, the digital twin system 40000 may extrapolate the make and model of the sensor.

[2049] In example embodiments, the manufacturer may provide or make available digital twins of their products (e.g., sensors, devices, machinery, equipment, raw materials, and the like). In these embodiments, the digital twin creation module 40064 may import the digital twins of one or more

products that are identified in the environment and may embed those digital twins in the digital twin of the environment. In example embodiments, embedding a digital twin within another digital twin may include creating a relationship between the embedded digital twin with the other digital twin. In these embodiments, the manufacturer of the digital twin may define the behaviors and/or properties of the respective products. For example, a digital twin of a machine may define the manner by which the machine operates, the inputs/outputs of the machine, and the like. In this way, the digital twin of the machine may reflect the operation of the machine given a set of inputs.

[2050] In example embodiments, a user may define one or more processes that occur in an environment. In these embodiments, the user may define the steps in the process, the machines/devices that perform each step in the process, the inputs to the process, and the outputs of the process.

[2051] In example embodiments, the digital twin creation module 40064 may create a graph database that defines the relationships between a set of digital twins. In these embodiments, the digital twin creation module 40064 may create nodes for the environment, systems and subsystems of the environment, devices in the environment, sensors in the environment, workers that work in the environment, processes that are performed in the environment, and the like. In example embodiments, the digital twin creation module 40064 may write the graph database representing a set of digital twins to the digital twin datastore 40016.

[2052] In example embodiments, the digital twin creation module 40064 may, for each node, include any data relating to the entity in the node representing the entity. For example, in defining a node representing an environment, the digital twin creation module 40064 may include the dimensions, boundaries, layout, pathways, and other relevant spatial data in the node. Furthermore, the digital twin creation module 40064 may define a coordinate space with respect to the environment. In the case that the digital twin may be rendered, the digital twin creation module 40064 may include a reference in the node to any shapes, meshes, splines, surfaces, and the like that may be used to render the environment. In representing a system, subsystem, device, or sensor, the digital twin creation module 40064 may create a node for the respective entity and may include any relevant data. For example, the digital twin creation module 40064 may create a node representing a machine in the environment. In this example, the digital twin creation module 40064 may include the dimensions, behaviors, properties, location, and/or any other suitable data relating to the machine in the node representing the machine. The digital twin creation module 40064 may connect nodes of related entities with an edge, thereby creating a relationship between the entities. In doing so, the created relationship between the entities may define the type of relationship characterized by the edge. In representing a process, the digital twin creation module 40064 may create a node for the entire process or may create a node for each step in the process. In some of these embodiments, the digital twin creation module 40064 may relate the process nodes to the nodes that represent the machinery/devices that perform the steps in the process. In embodiments where an edge connects the process step nodes to the machinery/device that performs the process step, the edge or one of the nodes may contain information that indicates the input to the step, the

output of the step, the amount of time the step takes, the nature of processing of inputs to produce outputs, a set of states or modes the process can undergo, and the like.

[2053] In example embodiments, the digital twin update module 40066 updates sets of digital twins based on a current status of one or more industrial entities. In example embodiments, the digital twin update module 40066 receives sensor data from a sensor system 40030 of an industrial environment and updates the status of the digital twin of the industrial environment and/or digital twins of any affected systems, subsystems, devices, workers, processes, or the like. As discussed, the digital twin I/O system 40004 may receive the sensor data in one or more sensor packets. The digital twin I/O system 40004 may provide the sensor data to the digital twin update module 40066 and may identify the environment from which the sensor packets were received and the sensor that provided the sensor packet. In response to the sensor data, the digital twin update module 40066 may update a state of one or more digital twins based on the sensor data. In some of these embodiments, the digital twin update module 40066 may update a record (e.g., a node in a graph database) corresponding to the sensor that provided the sensor data to reflect the current sensor data. In some scenarios, the digital twin update module 40066 may identify certain areas within the environment that are monitored by the sensor and may update a record (e.g., a node in a graph database) to reflect the current sensor data. For example, the digital twin update module 40066 may receive sensor data reflecting different vibrational characteristics of a machine and/or machine components. In this example, the digital twin update module 40066 may update the records representing the vibration sensors that provided the vibration sensor data and/or the records representing the machine and/or the machine components to reflect the vibration sensor data. In another example, in some scenarios, workers in an industrial environment (e.g., manufacturing facility, industrial storage facility, a mine, a drilling operation, or the like) may be required to wear wearable devices (e.g., smart watches, smart helmets, smart shoes, or the like). In these embodiments, the wearable devices may collect sensor data relating to the worker (e.g., location, movement, heartrate, respiration rate, body temperature, or the like) and/or the environment surrounding the worker and may communicate the collected sensor data to the digital twin system 40000 (e.g., via the real-time sensor connectivity 40014 such as a webhook) either directly or via an aggregation device of the sensor system. In response to receiving the sensor data from the wearable device of a worker, the digital twin update module 40066 may update a digital twin of a worker to reflect, for example, a location of the worker, a trajectory of the worker, a health status of the worker, or the like. In some of these embodiments, the digital twin update module 40066 may update a node representing a worker and/or an edge that connects the node representing the environment with the collected sensor data to reflect the current status of the worker.

[2054] In example embodiments, the digital twin update module 40066 may provide the sensor data from one or more sensors to the digital twin dynamic model system 40008, which may model a behavior of the environment and/or one or more industrial entities to extrapolate additional state data.

[2055] In example embodiments, the digital twin visualization module 40068 receives requests to view a visual digital twin or a portion thereof. In example embodiments, the request may indicate

the digital twin to be viewed (e.g., an environment identifier). In response, the digital twin visualization module 40068 may determine the requested digital twin and any other digital twins implicated by the request. For example, in requesting to view a digital twin of an environment, the digital twin visualization module 40068 may further identify the digital twins of any industrial entities within the environment. In example embodiments, the digital twin visualization module 40068 may identify the spatial relationships between the industrial entities and the environment based on, for example, the relationships defined in a graph database. In these embodiments, the digital twin visualization module 40068 can determine the relative location of embedded digital twins within the containing digital twin, relative locations of adjoining digital twins, and/or the transience of the relationship (e.g., is an object fixed to a point or does the object move). The digital twin visualization module 40068 may render the requested digital twins and any other implicated digital twin based on the identified relationships. In example embodiments, the digital twin visualization module 40068 may, for each digital twin, determine the surfaces of the digital twin. In example embodiments, the surfaces of a digital may be defined or referenced in a record corresponding to the digital twin, which may be provided by a user, determined from imported images, or defined by a manufacturer of an industrial entity. In the scenario that an object can take different poses or shapes (e.g., an industrial robot), the digital twin visualization module 40068 may determine a pose or shape of the object for the digital twin. The digital twin visualization module 40068 may embed the digital twins into the requested digital twin and may output the requested digital twin to a client application.

[2056] In example embodiments, the digital twin update module 40004 may provide the sensor data from one or more sensors to the digital twin dynamic system 40008, which may model a behavior of the environment and/or one or more industrial entities to extrapolate additional state data. For example, if an industrial storage facility includes temperature sensors 40032 at the four corners of a large space and each of the temperature sensors 40032 outputs a respective temperature reading corresponding to the ambient temperature surrounding the temperature sensor 40032, the digital twin dynamic system 40008 may determine temperatures in other unmonitored areas of the industrial storage facility. In this example, the digital twin dynamic system 40008 may output the determined temperatures to the digital twin update module 40004, which may update the digital twin of the environment to reflect the extrapolated temperatures. In these example embodiments, the determined temperatures may be used in any number of downstream applications. In example embodiments, the digital twin system 40004 may output the extrapolated temperatures (and the sensor-measured temperatures) to a monitoring system that classifies overheating conditions in the environment or improper temperatures. For example, the digital twin system 40004 may output an extrapolated temperature of a bearing, such that a determination of an overheated bearing may be indicative of a failure in rotating machinery. In another example, the digital twin system 40004 may output an extrapolated temperature of a brake pad, such that a determination of an overheated brake pad may be indicative of a brake failure. In another example, the digital twin system 40004 may output an extrapolated temperature of a food production facility, such that an improper temperature (e.g., below a minimum threshold or above an upper threshold) may lead to spoiling

perishables. In another example, the digital twin system 40004 may output an extrapolated temperature relating to a chemical process, such that an improper temperature (e.g., below a minimum threshold or above an upper threshold) may result in a failure of the chemical-based process. In another example, the digital twin system 40004 may output an extrapolated temperature of a cultivation facility, such that an improper temperature (e.g., below a minimum threshold or above an upper threshold) may lead to a crop failure. In another example, the digital twin system 40004 may output the extrapolated temperatures (and the sensor-measured temperatures) to a control system (e.g., HVAC controller) that adjusts the temperature within the environment based on the extrapolated and/or the sensor-measured temperatures.

[2057] In some of these embodiments, the request to view a digital twin may further indicate the type of view. As discussed, in example embodiments, digital twins may be depicted in a number of different view types. For example, an environment or device may be viewed in a "real-world" view that depicts the environment or device as they typically appear, in a "heat" view that depicts the environment or device in a manner that is indicative of a temperature of the environment or device, in a "vibration" view that depicts the machines and/or machine components in an industrial environment in a manner that is indicative of vibrational characteristics of the machines and/or machine components, in a "filtered" view that only displays certain types of objects within an environment or components of a device (such as objects that require attention resulting from, for example, recognition of a fault condition, an alert, an updated report, or other factor), an augmented view that overlays data on the digital twin, and/or any other suitable view types. In example embodiments, digital twins may be depicted in a number of different role-based view types. For example, a manufacturing facility device may be viewed in an "operator" view that depicts the facility in a manner that is suitable for a facility operator, a "C-Suite" view that depicts the facility in a manner that is suitable for executive-level managers, a "marketing" view that depicts the facility in a manner that is suitable for workers in sales and/or marketing roles, a "board" view that depicts the facility in a manner that is suitable for members of a corporate board, a "regulatory" view that depicts the facility in a manner that is suitable for regulatory managers, and a "human resources" view that depicts the facility in a manner that is suitable for human resources personnel. In response to a request that indicates a view type, the digital twin visualization module 40068 may retrieve the data for each digital twin that corresponds to the view type. For example, if a user has requested a vibration view of a factory floor, the digital twin visualization module 40068 may retrieve vibration data for the factory floor (which may include vibration measurements taken from different machines and/or machine components and/or vibration measurements that were extrapolated by the digital twin dynamic model system 40008 and/or simulated vibration data from digital twin simulation system 40006) as well as available vibration data for any industrial entities appearing on the factory floor. In this example, the digital twin visualization module 40068 may determine colors corresponding to each machine component on a factory floor that represent a vibration fault level state (e.g., red for alarm, orange for critical, yellow for suboptimal, and green for normal operation). The digital twin visualization module 40068 may then render the digital twins of the machine components within the environment based on the determined colors.

Additionally or alternatively, the digital twin visualization module 40068 may render the digital twins of the machine components within the environment with indicators having the determined colors. For instance, if the vibration fault level state of an inbound bearing of a motor is suboptimal and the outbound bearing of the motor is critical, the digital twin visualization module 40068 may render the digital twin of the inbound bearing having an indicator in a shade of yellow (e.g., suboptimal) and the outbound bearing having an indicator in a shade of orange (e.g., critical). It is noted that in example embodiments, the digital twin system 40000 may include an analytics system that determines the manner by which the digital twin visualization system presents information to a human user. For example, the analytics system may track outcomes relating to human interactions with real-world environments or objects in response to information presented in a visual digital twin. In example embodiments, the analytics system may apply cognitive models to determine the most effective manner to display visualized information (e.g., what colors to use to denote an alarm condition, what kind of movements or animations bring attention to an alarm condition, or the like) or audio information (what sounds to use to denote an alarm condition) based on the outcome data. In example embodiments, the analytics system may apply cognitive models to determine the most suitable manner to display visualized information based on the role of the user. In example embodiments, the visualization may include display of information related to the visualized digital twins, including graphical information, graphical information depicting vibration characteristics, graphical information depicting harmonic peaks, graphical information depicting peaks, vibration severity units data, vibration fault level state data, recommendations from cognitive intelligence system 40010, predictions from cognitive intelligence system 40010, probability of failure data, maintenance history data, time to failure data, cost of downtime data, probability of downtime data, cost of repair data, cost of machine replace data, probability of shutdown data, manufacturing KPIs, and the like.

[2058] In another example, a user may request a filtered view of a digital twin of a process, whereby the digital twin of the process only shows components (e.g., machine or equipment) that are involved in the process. In this example, the digital twin visualization module 40068 may retrieve a digital twin of the process, as well as any related digital twins (e.g., a digital twin of the environment and digital twins of any machinery or devices that impact the process). The digital twin visualization module 40068 may then render each of the digital twins (e.g., the environment and the relevant industrial entities) and then may perform the process on the rendered digital twins. It is noted that as a process may be performed over a period of time and may include moving items and/or parts, the digital twin visualization module 40068 may generate a series of sequential frames that demonstrate the process. In this scenario, the movements of the machines and/or devices implicated by the process may be determined according to the behaviors defined in the respective digital twins of the machines and/or devices.

[2059] As discussed, the digital twin visualization module 40068 may output the requested digital twin to a client application 40070. In example embodiments, the client application 40070 is a virtual reality application, whereby the requested digital twin is displayed on a virtual reality headset. In example embodiments, the client application 40070 is an augmented reality application,

whereby the requested digital twin is depicted in an AR-enabled device. In these embodiments, the requested digital twin may be filtered such that visual elements and/or text are overlaid on the display of the AR-enabled device.

[2060] It is noted that while a graph database is discussed, the digital twin system 40000 may employ other suitable data structures to store information relating to a set of digital twins. In these embodiments, the data structures, and any related storage system, may be implemented such that the data structures provide for some degree of feedback loops and/or recursion when representing iteration of flows.

[2061] FIG. 215 illustrates an example of a digital twin I/O system 40004 that interfaces with the environment 40020, the digital twin system 40000, and/or components thereof to provide bi-directional transfer of data between coupled components according to some embodiments of the present disclosure.

[2062] In example embodiments, the transferred data includes signals (e.g., request signals, command signals, response signals, etc.) between connected components, which may include software components, hardware components, physical devices, virtualized devices, simulated devices, combinations thereof, and the like. The signals may define material properties (e.g., physical quantities of temperature, pressure, humidity, density, viscosity, etc.), measured values (e.g., contemporaneous or stored values acquired by the device or system), device properties (e.g., device ID or properties of the device's design specifications, materials, measurement capabilities, dimensions, absolute position, relative position, combinations thereof, and the like), set points (e.g., targets for material properties, device properties, system properties, combinations thereof, and the like), and/or critical points (e.g., threshold values such as minimum or maximum values for material properties, device properties, system properties, etc.). The signals may be received from systems or devices that acquire (e.g., directly measure or generate) or otherwise obtain (e.g., receive, calculate, look-up, filter, etc.) the data, and may be communicated to or from the digital twin I/O system 40004 at predetermined times or in response to a request (e.g., polling) from the digital twin I/O system 40004. The communications may occur through direct or indirect connections (e.g., via intermediate modules within a circuit and/or intermediate devices between the connected components). The values may correspond to real-world elements 40302r (e.g., an input or output for a tangible vibration sensor) or virtual elements 40302v (e.g., an input or output for a digital twin 40302d and/or a simulated element 40302s that provide vibration data).

[2063] In example embodiments, the real-world elements 40302r may be elements within the industrial environment 40020. The real-world elements 40302r may include, for example, non-networked elements 40022, the devices 40024 (smart or non-smart), sensors 40026, and humans 40028. The real-world elements 40302r may be process or non-process equipment within the industrial environments 40020. For example, process equipment may include motors, pumps, mills, fans, painters, welders, smelters, etc., and non-process equipment may include personal protective equipment, safety equipment, emergency stations or devices (e.g., safety showers, eyewash stations, fire extinguishers, sprinkler systems, etc.), warehouse features (e.g., walls, floor layout, etc.), obstacles (e.g., persons or other items within the environment 40020, etc.), etc.

[2064] In example embodiments, the virtual elements 40302v may be digital representations of or that correspond to contemporaneously existing real-world elements 40302r. Additionally or alternatively, the virtual elements 40302v may be digital representations of or that correspond to real-world elements 40302r that may be available for later addition and implementation into the environment 40020. The virtual elements may include, for example, simulated elements 40302s and/or digital twins 40302d. In example embodiments, the simulated elements 40302s may be digital representations of real-world elements 40302s that are not present within the industrial environment 40020. The simulated elements 40302s may mimic desired physical properties which may be later integrated within the environment 40020 as real-world elements 40302r (e.g., a "black box" that mimics the dimensions of a real-world elements 40302r). The simulated elements 40302s may include digital twins of existing objects (e.g., a single simulated element 40302s may include one or more digital twins 40302d for existing sensors). Information related to the simulated elements 40302s may be obtained, for example, by evaluating behavior of corresponding real-world elements 40302r using mathematical models or algorithms, from libraries that define information and behavior of the simulated elements 40302s (e.g., physics libraries, chemistry libraries, or the like).

[2065] In example embodiments, the digital twin 40302d may be a digital representation of one or more real-world elements 40302r. The digital twins 40302d are configured to mimic, copy, and/or model behaviors and responses of the real-world elements 40302r in response to inputs, outputs, and/or conditions of the surrounding or ambient environment. Data related to physical properties and responses of the real-world elements 40302r may be obtained, for example, via user input, sensor input, and/or physical modeling (e.g., thermodynamic models, electrodynamic models, mechanodynamic models, etc.). Information for the digital twin 40302d may correspond to and be obtained from the one or more real-world elements 40302r corresponding to the digital twin 40302d. For example, in example embodiments, the digital twin 40302d may correspond to one real-world element 40302r that is a fixed digital vibration sensor 40036 on a machine component, and vibration data for the digital twin 40302d may be obtained by polling or fetching vibration data measured by the fixed digital vibration sensor on the machine component. In a further example, the digital twin 40302d may correspond to a plurality of real-world elements 40302r such that each of the elements can be a fixed digital vibration sensor on a machine component, and vibration data for the digital twin 40302d may be obtained by polling or fetching vibration data measured by each of the fixed digital vibration sensors on the plurality of real-world elements 40302r. Additionally or alternatively, vibration data of a first digital twin 40302d may be obtained by fetching vibration data of a second digital twin 40302d that is embedded within the first digital twin 40302d, and vibration data for the first digital twin 40302d may include or be derived from vibration data for the second digital twin 40302d. For example, the first digital twin may be a digital twin 40302d of an environment 40020 (alternatively referred to as an "environmental digital twin") and the second digital twin 40302d may be a digital twin 40302d corresponding to a vibration sensor disposed within the environment 40020 such that the vibration data for the first digital twin 40302d is

obtained from or calculated based on data including the vibration data for the second digital twin 40302d.

[2066] In example embodiments, the digital twin system 40000 monitors properties of the real-world elements 40302r using the sensors 40026 within a respective environment 40020 that is or may be represented by a digital twin 40302d and/or outputs of models for one or more simulated elements 40302s. In example embodiments, the digital twin system 40000 may minimize network congestion while maintaining effective monitoring of processes by extending polling intervals and/or minimizing data transfer for sensors corresponding that correspond to affected real-world elements 40302r and performing simulations (e.g., via the digital-twin simulation system 106) during the extended interval using data that was obtained from other sources (e.g., sensors that are physically proximate to or have an effect on the affected real-world elements 40302r). Additionally or alternatively, error checking may be performed by comparing the collected sensor data with data obtained from the digital-twin simulation system 106. For example, consistent deviations or fluctuations between sensor data obtained from the real-world element 40302r and the simulated element 40302s may indicate malfunction of the respective sensor or another fault condition.

[2067] In example embodiments, the digital twin system 40000 may optimize features of the environment through use of one or more simulated elements 40302s. For example, the digital twin system 40000 may evaluate effects of the simulated elements 40302s within a digital twin of an environment to quickly and efficiently determine costs and/or benefits flowing from inclusion, exclusion, or substitution of real-world elements 40302r within the environment 40020. The costs and benefits may include, for example, increased machinery costs (e.g., capital investment and maintenance), increased efficiency (e.g., process optimization to reduce waste or increase throughput), decreased or altered footprint within the environment 40020, extension or optimization of useful lifespans, minimization of component faults, minimization of component downtime, etc.

[2068] In example embodiments, the digital twin I/O system 40004 may include one or more software modules that are executed by one or more controllers of one or more devices (e.g., server devices, user devices, and/or distributed devices) to affect the described functions. The digital twin I/O system 40004 may include, for example, an input module 400304, an output module 400306, and an adapter module 400308.

[2069] In example embodiments, the input module 400304 may obtain or import data from data sources in communication with the digital twin I/O system 40004, such as the sensor system 40030 and the digital twin simulation system 40006. The data may be immediately used by or stored within the digital twin system 40000. The imported data may be ingested from data streams, data batches, in response to a triggering event, combinations thereof, and the like. The input module 400304 may receive data in a format that is suitable to transfer, read, and/or write information within the digital twin system 40000.

[2070] In example embodiments, the output module 400306 may output or export data to other system components (e.g., the digital twin datastore 40016, the digital twin simulation system 40006, the cognitive intelligence system 40010, etc.), devices 40024, and/or the client application

40070. The data may be output in data streams, data batches, in response to a triggering event (e.g., a request), combinations thereof, and the like. The output module 400306 may output data in a format that is suitable to be used or stored by the target element (e.g., one protocol for output to the client application and another protocol for the digital twin datastore 40016).

[2071] In example embodiments, the adapter module 400308 may process and/or convert data between the input module 400304 and the output module 400306. In example embodiments, the adapter module 400308 may convert and/or route data automatically (e.g., based on data type) or in response to a received request (e.g., in response to information within the data).

[2072] In example embodiments, the digital twin system 40000 may represent a set of industrial workpiece elements in a digital twin, and the digital twin simulation system 40006 simulates a set of physical interactions of a worker with the workpiece elements. The simulated physical interactions may include, for example, workpiece movements (e.g., a worker carrying the workpiece between locations), placement of the workpiece (e.g., a worker mounting or aligning the workpiece for further processing), machine actuation (e.g., machine-bending sheet metal in response to placement of the workers hands and/or feet on designated triggers), manual workpiece alterations (e.g., the worker painting, welding, and/or removing material from the workpiece by hand), and the like.

[2073] In example embodiments, the digital twin simulation system 40006 may determine process outcomes for the simulated physical interactions accounting for simulated human factors. For example, variations in workpiece throughput may be modeled by the digital twin system 40000 including, for example, worker response times to events, worker fatigue, discontinuity within worker actions (e.g., natural variations in human-movement speed, differing positioning times, etc.), effects of discontinuities on downstream processes, and the like. In example embodiments, individualized worker interactions may be modeled using historical data that is collected, acquired, and/or stored by the digital twin system 40000. The simulation may begin based on estimated amounts (e.g., worker age, industry averages, workplace expectations, etc.). The simulation may also individualize data for each worker (e.g., comparing estimated amounts to collected worker-specific outcomes).

[2074] In example embodiments, information relating to workers (e.g., fatigue rates, efficiency rates, and the like) may be determined by analyzing performance of specific workers over time and modeling said performance.

[2075] In example embodiments, the digital twin system 40000 includes a plurality of proximity sensors within the sensor array 40030. The proximity sensors are or may be configured to detect elements of the environment 40020 that are within a predetermined area. For example, proximity sensors may include electromagnetic sensors, light sensors, and/or acoustic sensors.

[2076] The electromagnetic sensors are or may be configured to sense objects or interactions via one or more electromagnetic fields (e.g., emitted electromagnetic radiation or received electromagnetic radiation). In example embodiments, the electromagnetic sensors include inductive sensors (e.g., radio-frequency identification sensors), capacitive sensors (e.g., contact and contactless capacitive sensors), combinations thereof, and the like.

[2077] The light sensors are or may be configured to sense objects or interactions via electromagnetic radiation in, for example, the far-infrared, near-infrared, optical, and/or ultraviolet spectra. In example embodiments, the light sensors may include image sensors (e.g., charge-coupled devices and CMOS active-pixel sensors), photoelectric sensors (e.g., through-beam sensors, retroreflective sensors, and diffuse sensors), combinations thereof, and the like. Further, the light sensors may be implemented as part of a system or subsystem, such as a light detection and ranging ("LIDAR") sensor.

[2078] The acoustic sensors are or may be configured to sense objects or interactions via sound waves that are emitted and/or received by the acoustic sensors. In example embodiments, the acoustic sensors may include infrasonic, sonic, and/or ultrasonic sensors. Further, the acoustic sensors may be grouped as part of a system or subsystem, such as a sound navigation and ranging ("SONAR") sensor.

[2079] In example embodiments, the digital twin system 40000 stores and collects data from a set of proximity sensors within the environment 40020 or portions thereof. The collected data may be stored, for example, in the digital twin datastore 40016 for use by components the digital twin system 40000 and/or visualization by a user. Such use and/or visualization may occur contemporaneously with or after collection of the data (e.g., during later analysis and/or optimization of processes).

[2080] In example embodiments, data collection may occur in response to a triggering condition. These triggering conditions may include, for example, expiration of a static or a dynamic predetermined interval, obtaining a value short of or in excess of a static or dynamic value, receiving an automatically generated request or instruction from the digital twin system 40000 or components thereof, interaction of an element with the respective sensor or sensors (e.g., in response to a worker or machine breaking a beam or coming within a predetermined distance from the proximity sensor), interaction of a user with a digital twin (e.g., selection of an environmental digital twin, a sensor array digital twin, or a sensor digital twin), combinations thereof, and the like.

[2081] In example embodiments, the digital twin system 40000 collects and/or stores RFID data in response to interaction of a worker with a real-world element 40302r. For example, in response to a worker interaction with a real-world environment, the digital twin will collect and/or store RFID data from RFID sensors within or associated with the corresponding environment 40020. Additionally or alternatively, worker interaction with a sensor-array digital twin will collect and/or store RFID data from RFID sensors within or associated with the corresponding sensor array. Similarly, worker interaction with a sensor digital twin will collect and/or store RFID data from the corresponding sensor. The RFID data may include suitable data attainable by RFID sensors such as proximate RFID tags, RFID tag position, authorized RFID tags, unauthorized RFID tags, unrecognized RFID tags, RFID type (e.g., active or passive), error codes, combinations thereof, and the like.

[2082] In example embodiments, the digital twin system 40000 may further embed outputs from one or more devices within a corresponding digital twin. In example embodiments, the digital twin

system 40000 embeds output from a set of individual-associated devices into an industrial digital twin. For example, the digital twin I/O system 40004 may receive information output from one or more wearable devices 40054 or mobile devices (not shown) associated with an individual within an industrial environment. The wearable devices may include image capture devices (e.g., body cameras or augmented-reality headwear), navigation devices (e.g., GPS devices, inertial guidance systems), motion trackers, acoustic capture devices (e.g., microphones), radiation detectors, combinations thereof, and the like.

[2083] In example embodiments, upon receiving the output information, the digital twin I/O system 40004 routes the information to the digital twin creation module 40064 to check and/or update the environment digital twin and/or associated digital twins within the environment (e.g., a digital twin of a worker, machine, or robot position at a given time). Further, the digital twin system 40000 may use the embedded output to determine characteristics of the environment 40020.

[2084] In example embodiments, the digital twin system 40000 embeds output from a LIDAR point cloud system into an industrial digital twin. For example, the digital twin I/O system 40004 may receive information output from one or more Lidar devices 40038 within an industrial environment. The Lidar devices 40038 is configured to provide a plurality of points having associated position data (e.g., coordinates in absolute or relative x, y, and z values). Each of the plurality of points may include further LIDAR attributes, such as intensity, return number, total returns, laser color data, return color data, scan angle, scan direction, etc. The Lidar devices 40038 may provide a point cloud that includes the plurality of points to the digital twin system 40000 via, for example, the digital twin I/O system 40004. Additionally or alternatively, the digital twin system 40000 may receive a stream of points and assemble the stream into a point cloud, or may receive a point cloud and assemble the received point cloud with existing point cloud data, map data, or three dimensional (3D)-model data.

[2085] In example embodiments, upon receiving the output information, the digital twin I/O system 40004 routes the point cloud information to the digital twin creation module 40064 to check and/or update the environment digital twin and/or associated digital twins within the environment (e.g., a digital twin of a worker, machine, or robot position at a given time). In example embodiments, the digital twin system 40000 is further configured to determine closed-shape objects within the received LIDAR data. For example, the digital twin system 40000 may group a plurality of points within the point cloud as an object and, if necessary, estimate obstructed faces of objects (e.g., a face of the object contacting or adjacent a floor or a face of the object contacting or adjacent another object such as another piece of equipment). The system may use such closed-shape objects to narrow the search space for digital twins and thereby increase efficiency of matching algorithms (e.g., a shape-matching algorithm).

[2086] In example embodiments, the digital twin system 40000 embeds output from a simultaneous location and mapping ("SLAM") system in an environmental digital twin. For example, the digital twin I/O system 40004 may receive information output from the SLAM system, such as Slam sensor 40062, and embed the received information within an environment digital twin corresponding to the location determined by the SLAM system. In example

embodiments, upon receiving the output information from the SLAM system, the digital twin I/O system 40004 routes the information to the digital twin creation module 40064 to check and/or update the environment digital twin and/or associated digital twins within the environment (e.g., a digital twin of a workpiece, furniture, movable object, or autonomous object). Such updating provides digital twins of non-connected elements (e.g., furnishings or persons) automatically and without need of user interaction with the digital twin system 40000.

[2087] In example embodiments, the digital twin system 40000 can leverage known digital twins to reduce computational requirements for the Slam sensor 40062 by using suboptimal map-building algorithms. For example, the suboptimal map-building algorithms may allow for a higher uncertainty tolerance using simple bounded-region representations and identifying possible digital twins. Additionally or alternatively, the digital twin system 40000 may use a bounded-region representation to limit the number of digital twins, analyze the group of potential twins for distinguishing features, then perform higher precision analysis for the distinguishing features to identify and/or eliminate categories of, groups of, or individual digital twins and, in the event that no matching digital twin is found, perform a precision scan of only the remaining areas to be scanned.

[2088] In example embodiments, the digital twin system 40000 may further reduce compute required to build a location map by leveraging data captured from other sensors within the environment (e.g., captured images or video, radio images, etc.) to perform an initial map-building process (e.g., a simple bounded-region map or other suitable photogrammetry methods), associate digital twins of known environmental objects with features of the simple bounded-region map to refine the simple bounded-region map, and perform more precise scans of the remaining simple bounded regions to further refine the map. In example embodiments, the digital twin system 40000 may detect objects within received mapping information and, for each detected object, determine whether the detected object corresponds to an existing digital twin of a real-world-element. In response to determining that the detected object does not correspond to an existing real-world-element digital twin, the digital twin system 40000 may use, for example, the digital twin creation module 40064 to generate a new digital twin corresponding to the detected object (e.g., a detected-object digital twin) and add the detected-object digital twin to the real-world-element digital twins within the digital twin datastore. Additionally or alternatively, in response to determining that the detected object corresponds to an existing real-world-element digital twin, the digital twin system 40000 may update the real-world-element digital twin to include new information detected by the simultaneous location and mapping sensor, if any.

[2089] In example embodiments, the digital twin system 40000 represents locations of autonomously or remotely moveable elements and attributes thereof within an industrial digital twin. Such movable elements may include, for example, workers, persons, vehicles, autonomous vehicles, robots, etc. The locations of the moveable elements may be updated in response to a triggering condition. Such triggering conditions may include, for example, expiration of a static or a dynamic predetermined interval, receiving an automatically generated request or instruction from the digital twin system 40000 or components thereof, interaction of an element with a respective

sensor or sensors (e.g., in response to a worker or machine breaking a beam or coming within a predetermined distance from a proximity sensor), interaction of a user with a digital twin (e.g., selection of an environmental digital twin, a sensor array digital twin, or a sensor digital twin), combinations thereof, and the like.

[2090] In example embodiments, the time intervals may be based on probability of the respective movable element having moved within a time period. For example, the time interval for updating a worker location may be relatively shorter for workers expected to move frequently (e.g., a worker tasked with lifting and carrying objects within and through the environment 40020) and relatively longer for workers expected to move infrequently (e.g., a worker tasked with monitoring a process stream). Additionally or alternatively, the time interval may be dynamically adjusted based on applicable conditions, such as increasing the time interval when no movable elements are detected, decreasing the time interval as or when the number of moveable elements within an environment increases (e.g., increasing number of workers and worker interactions), increasing the time interval during periods of reduced environmental activity (e.g., breaks such as lunch), decreasing the time interval during periods of abnormal environmental activity (e.g., tours, inspections, or maintenance), decreasing the time interval when unexpected or uncharacteristic movement is detected (e.g., frequent movement by a typically sedentary element or coordinated movement, for example, of workers approaching an exit or moving cooperatively to carry a large object), combinations thereof, and the like. Further, the time interval may also include additional, semi-random acquisitions. For example, occasional mid-interval locations may be acquired by the digital twin system 40000 to reinforce or evaluate the efficacy of the particular time interval.

[2091] In example embodiments, the digital twin system 40000 may analyze data received from the digital twin I/O system 40004 to refine, remove, or add conditions. For example, the digital twin system 40000 may optimize data collection times for movable elements that are updated more frequently than needed (e.g., multiple consecutive received positions being identical or within a predetermined margin of error).

[2092] In example embodiments, the digital twin system 40000 may receive, identify, and/or store a set of states 40040a-n related to the environment 40020. The states 40040a-n may be, for example, data structures that include a plurality of attributes 40404a-n and a set of identifying criteria 40406a-n to uniquely identify each respective state 40040a-n. In example embodiments, the states 40040a-n may correspond to states where it is desirable for the digital twin system 40000 to set or alter conditions of real-world elements 40302r and/or the environment 40020 (e.g., increase/decrease monitoring intervals, alter operating conditions, etc.).

[2093] In example embodiments, the set of states 40040a-n may further include, for example, minimum monitored attributes for each state 40040a-n, the set of identifying criteria 40406a-n for each state 40040a-n, and/or actions available to be taken or recommended to be taken in response to each state 40040a-n. Such information may be stored by, for example, the digital twin datastore 40016 or another datastore. The states 40040a-n or portions thereof may be provided to, determined by, or altered by the digital twin system 40000. Further, the set of states 40040a-n may include data from disparate sources. For example, details to identify and/or respond to occurrence of a first

state may be provided to the digital twin system 40000 via user input, details to identify and/or respond to occurrence of a second state may be provided to the digital twin system 40000 via an external system, details to identify and/or respond to occurrence of a third state may be determined by the digital twin system 40000 (e.g., via simulations or analysis of process data), and details to identify and/or respond to occurrence of a fourth state may be stored by the digital twin system 40000 and altered as desired (e.g., in response to simulated occurrence of the state or analysis of data collected during an occurrence of and response to the state).

[2094] In example embodiments, the plurality of attributes 40404a-n includes at least the attributes 40404a-n needed to identify the respective state 40040a-n. The plurality of attributes 40404a-n may further include additional attributes that are or may be monitored in determining the respective state 40040a-n, but are not needed to identify the respective state 40040a-n. For example, the plurality of attributes 40404a-n for a first state may include relevant information such as rotational speed, fuel level, energy input, linear speed, acceleration, temperature, strain, torque, volume, weight, etc.

[2095] The set of identifying criteria 40406a-n may include information for each of the set of attributes 40404a-n to uniquely identify the respective state. The identifying criteria 40406a-n may include, for example, rules, thresholds, limits, ranges, logical values, conditions, comparisons, combinations thereof, and the like.

[2096] The change in operating conditions or monitoring may be any suitable change. For example, after identifying occurrence of a respective state 40406a-n, the digital twin system 40000 may increase or decrease monitoring intervals for a device (e.g., decreasing monitoring intervals in response to a measured parameter differing from nominal operation) without altering operation of the device. Additionally or alternatively, the digital twin system 40000 may alter operation of the device (e.g., reduce speed or power input) without altering monitoring of the device. In further embodiments, the digital twin system 40000 may alter operation of the device (e.g., reduce speed or power input) and alter monitoring intervals for the device (e.g., decreasing monitoring intervals).

[2097] FIG. 216 illustrates an example set of identified states 40040a-n related to industrial environments that the digital twin system 40000 may identify and/or store for access by intelligent systems (e.g., the cognitive intelligence system 40010) or users of the digital twin system 40000, according to some embodiments of the present disclosure. The states 40040a-n may include operational states (e.g., suboptimal, normal, optimal, critical, or alarm operation of one or more components), excess or shortage states (e.g., supply-side or output-side quantities), combinations thereof, and the like.

[2098] In example embodiments, the digital twin system 40000 may monitor attributes 40404a-n of real-world elements 40302r and/or digital twins 40302d to determine the respective state 40040a-n. The attributes 40404a-n may be, for example, operating conditions, set points, critical points, status indicators, other sensed information, combinations thereof, and the like. For example, the attributes 40404a-n may include power input 40404a, operational speed 40404b, critical speed 40404c, and operational temperature 40404d of the monitored elements. While the illustrated example illustrates uniform monitored attributes, the monitored attributes may differ by target

device (e.g., the digital twin system 40000 would not monitor rotational speed for an object with no rotatable components).

[2099] Each of the states 40040a-n includes a set of identifying criteria 40406a-n meeting particular criteria that are unique among the group of monitored states 40040a-n. The digital twin system 40000 may identify the overspeed state 40040a, for example, in response to the monitored attributes 40404a-n meeting a first set of identifying criteria 40406a (e.g., operational speed 40404b being higher than the critical speed 40404c, while the operational temperature 40404d is nominal).

[2100] In response to determining that one or more states 40040a-n exists or has occurred, the digital twin system 40000 may update triggering conditions for one or more monitoring protocols, issue an alert or notification, or trigger actions of subcomponents of the digital twin system 40000. For example, subcomponents of the digital twin system 40000 may take actions to mitigate and/or evaluate impacts of the detected states 40040a-n. When attempting to take actions to mitigate impacts of the detected states 40040a-n on real-world elements 40302r, the digital twin system 40000 may determine whether instructions exist (e.g., are stored in the digital twin datastore 40016) or should be developed (e.g., developed via simulation and cognitive intelligence or via user or worker input). Further, the digital twin system 40000 may evaluate impacts of the detected states 40040a-n, for example, concurrently with the mitigation actions or in response to determining that the digital twin system 40000 has no stored mitigation instructions for the detected states 40040a-n.

[2101] In example embodiments, the digital twin system 40000 employs the digital twin simulation system 40006 to simulate one or more impacts, such as immediate, upstream, downstream, and/or continuing effects, of recognized states. The digital twin simulation system 40006 may collect and/or be provided with values relevant to the evaluated states 40040a-n. In simulating the impact of the one or more states 40040a-n, the digital twin simulation system 40006 may recursively evaluate performance characteristics of affected digital twins 40302d until convergence is achieved. The digital twin simulation system 40006 may work, for example, in tandem with the cognitive intelligence system 40010 to determine response actions to alleviate, mitigate, inhibit, and/or prevent occurrence of the one or more states 40040a-n. For example, the digital twin simulation system 40006 may recursively simulate impacts of the one or more states 40040a-n until achieving a desired fit (e.g., convergence is achieved), provide the simulated values to the cognitive intelligence system 40010 for evaluation and determination of potential actions, receive the potential actions, evaluate impacts of each of the potential actions for a respective desired fit (e.g., cost functions for minimizing production disturbance, preserving critical components, minimizing maintenance and/or downtime, optimizing system, worker, user, or personal safety, etc.).

[2102] In example embodiments, the digital twin simulation system 40006 and the cognitive intelligence system 40010 may repeatedly share and update the simulated values and response actions for each desired outcome until desired conditions are met (e.g., convergence for each evaluated cost function for each evaluated action). The digital twin system 40000 may store the results in the digital twin datastore 40016 for use in response to determining that one or more states

40040a-n has occurred. Additionally, simulations and evaluations by the digital twin simulation system 40006 and/or the cognitive intelligence system 40010 may occur in response to occurrence or detection of the event.

[2103] In example embodiments, simulations and evaluations are triggered only when associated actions are not present within the digital twin system 40000. In further embodiments, simulations and evaluations are performed concurrently with use of stored actions to evaluate the efficacy or effectiveness of the actions in real time and/or evaluate whether further actions should be employed or whether unrecognized states may have occurred. In example embodiments, the cognitive intelligence system 40010 may also be provided with notifications of instances of undesired actions with or without data on the undesired aspects or results of such actions to optimize later evaluations.

[2104] In example embodiments, the digital twin system 40000 evaluates and/or represents the impact of machine downtime within a digital twin of a manufacturing facility. For example, the digital twin system 40000 may employ the digital twin simulation system 40006 to simulate the immediate, upstream, downstream, and/or continuing effects of a machine downtime state 40040b. The digital twin simulation system 40006 may collect or be provided with performance-related values such as optimal, suboptimal, and minimum performance requirements for elements (e.g., real-world elements 40302r and/or nested digital twins 40302d) within the affected digital twins 40302d, and/or characteristics thereof that are available to the affected digital twins 40302d, nested digital twins 40302d, redundant systems within the affected digital twins 40302d, combinations thereof, and the like.

[2105] In example embodiments, the digital twin system 40000 is configured to: simulate one or more operating parameters for the real-world elements in response to the industrial environment being supplied with given characteristics using the real-world-element digital twins; calculate a mitigating action to be taken by one or more of the real-world elements in response to being supplied with the contemporaneous characteristics; and actuate, in response to detecting the contemporaneous characteristics, the mitigating action. The calculation may be performed in response to detecting contemporaneous characteristics or operating parameters falling outside of respective design parameters or may be determined via a simulation prior to detection of such characteristics.

[2106] Additionally or alternatively, the digital twin system 40000 may provide alerts to one or more users or system elements in response to detecting states.

[2107] In example embodiments, the digital twin system 40000 includes power source characteristics of an industrial environment in a digital twin. The power source characteristics may include, for example, potential power sources, available power from individual lines or the grid, battery-based devices that can share power with other elements of the environment, back-up power systems, as well as environmental power sources (e.g., heat sources that can be utilized and converted to power). The power source characteristics may further include delivered-power information, such as delivered power factor, power quality, utility frequency, circuit frequency, phase shifts (e.g., capacitance and inductance differences in power routing), time lag for

switchover, distribution lag (e.g., if devices or circuits require an amount of energy or reaching steady state prior to actuation), combinations thereof, and the like.

[2108] In example embodiments, the mitigating actions may include, for example, stopping power-consuming elements within the environment, reducing power supplied to one or more devices within the environment, providing power from an alternative power source external to the environment, allocating power from power storage devices within the environment, combinations thereof, and the like. The batteries and/or capacitors within the environment may be stand-alone elements (e.g., a battery bank or a capacitor bank) or integrated within elements of the environment (e.g., a battery pack within an electric vehicle or elements that have battery backups). Further, the mitigation actions may be performed iteratively such that additional actions may be taken in response to continuing power loss state. For example, the digital twin system 40000 may switch the environment to power supplied by a battery bank and stop a first set of power-consuming elements in response to detecting a power-loss state 40402b. In the event that the power-loss state 40402b continues through a particular triggering event, the digital twin system 40000 may take further actions, such as further reducing power consumption of the environment by stopping a second set of power consuming devices and/or reducing operation of a third set of power consuming devices. The mitigating actions may further include, for example, actuating one of an inductive circuit or a capacitive circuit operatively coupled between the power source and the real-world elements to optimize power supplied to real-world elements within the industrial environment.

[2109] In example embodiments, the triggering events may include, for example, the stored energy within the battery bank falling below a predetermined level, the digital twin system 40000 receiving notification that the power-loss state 40402b is expected to continue for a certain duration, the digital twin system 40000 determining that the power-loss state 40402b is expected to continue for a certain duration, combinations thereof, and the like.

[2110] Additionally or alternatively, the digital twin system 40000 may provide alerts to one or more users or system elements in response to detecting states, such as a power-loss state 40402b. For example, actions taken by the digital twin system 40000 may be able to prevent any noticeable impact on the environment from a power-loss state 40402b, so the digital twin system 40000 may provide an alert to users of the digital twin system 40000. The alert may be a notification of an occurrence of the power-loss state 40402b, an indication of data corresponding to the power-loss state 40402b (e.g., reliability statistics), instructions for reducing impact of future events (e.g., switching power sources in response to the power reliability dropping below a predetermined amount), instructions on effect of the particular power-loss state 40402b on the environment (e.g., altered maintenance schedule or devices that performed unexpectedly during the power-loss state 40402b), combinations thereof, and the like.

[2111] In example embodiments, the digital twin system 40000 may increase longevity of power-backup systems within the environment based on simulations performed by the simulation system. For example, determining the probability that such backup systems will be employed within a timeframe allows the backup systems to be maintained at reduced capacity. The probability

calculation may employ, for example, weather forecast data, contemporaneous weather data, historical data collected by the digital twin system 40000, simulation data based on data collected by the sensor array (e.g., unexpected power fluctuations indicative of an impending mechanical event), combinations thereof, and the like.

[2112] In example embodiments, a backup battery system is maintained at an optimum level below maximum capacity to thereby increase battery longevity while providing adequate backup capacity and minimizing total storage of the backup system. For example, the battery bank may be energized to about 80% of capacity and maintained at that level until the probability of a power outage (as determined by the digital twin system 40000) exceeds a predetermined threshold (e.g., 50% chance) within a given window (e.g., the time it takes to charge the backup system to capacity). In response to the probability exceeding the predetermined amount, the digital twin system 40000 may initiate charging of the batteries to full capacity. The digital twin system 40000 may maintain the charge at capacity until discharge is required by an outage or until the probability of an outage falls below another predetermined threshold (e.g., below 10%) within a given window. In response to the probability of an outage falling below the predetermined threshold while the backup system is above optimum charge, the digital twin system 40000 may selectively discharge the backup system to return to the optimum level or a desired level to promote battery health and longevity.

[2113] Additionally or alternatively, the digital twin system 40000 may leverage the probability calculation to minimize cost of the stored power. For example, in a stored-electricity backup system such as a battery bank or capacitor bank, the digital twin system 40000 may delay charging of the backup until a lower price of electricity is available (e.g., off-peak hours, wholesale price drops to a particular amount, solar or other renewable energy is available, etc.). Further, in a generating backup system (e.g., fuel-powered generators) that receives fuel from on-site storage tanks, the digital twin system 40000 may delay purchase of additional fuel until fuel prices meet a desired amount or the probability of an outage before delivery exceeds a predetermined threshold (e.g., delivery takes one week from order, and the digital twin system 40000 determines a probability of an outage due to a weather event proximate that lead time).

[2114] In example embodiments, the digital twin system 40000 evaluates and/or represents the impact of a network connectivity outage in a digital twin of a real-world network. For example, the digital twin system 40000 may employ the digital twin simulation system 40006 to simulate the immediate, upstream, downstream, and/or continuing effects of a network-constrained state. The network-constrained states may include, for example, connection loss or constraint (e.g., bandwidth loss, network congestion or bandwidth exhaustion, and latency increases), interference (e.g., intermittent connectivity, packet drops, and increased transfer overhead), signal strength reduction, data collisions, address exhaustion, combinations thereof, and the like.

[2115] In example embodiments, the digital twin simulation system 40006 may collect or be provided with network-related values such as optimal, suboptimal, and minimum bandwidth and/or quality of service requirements for real-world elements 40302r within or attached to the network, potential data transfer routes through the network, alternate connectivity capability of real-world elements 40302r within the network, effect of connectivity loss on real-world elements 40302r,

bandwidth reduction or latency increases within the network, redundant systems within the affected networks, combinations thereof, and the like. In example embodiments, the digital twin simulation system 40006 may simulate various network-constrained states by utilizing digital twins of the network or components thereof, such as simulating loss of one or more components within the environment, loss of connectivity between components, loss of communication between the environment and the WAN, bottlenecks, human interactions with the network connectivity components, bandwidth or connectivity changes from external events (e.g., rain, temperature, electromagnetic interference, increased transmissivity at night, etc.), increased signaling through the system (e.g., in response to one or more devices within the environment increasing polling or increasing sent values), combinations thereof, and the like. The digital twin simulation system 40006 may store such simulations within, for example, the digital twin datastore 40016 for later use.

[2116] In the context of a communication network, mitigating actions may include, for example, establishing a failover connection, establishing an ad-hoc network connection capable of routing data around affected devices, reducing data from one or more devices, increasing real-world elements 40302r capable of data transfer therethrough (e.g., increasing access points), allocate bandwidth of one or more WAN-attachable devices, combinations thereof, and the like. Data from one or more devices may be reduced, for example, by reducing polling intervals from low-priority or redundant devices, stopping data transfer from potentially redundant devices, pushing data processing toward the edge to reduce network throughput of raw data, etc. Bandwidth from the WAN-attachable devices may be allocated to serve affected portions of the network. As used herein, "WAN-attachable devices" are devices that can have direct connections to devices outside of the environment (e.g., to cellular towers or independent internet connections). For example, the wearable devices 40054 may include a Wi-Fi transmitter and receiver as well as a cellular transmitter and receiver capable of sending data via a cellular network. The digital twin system 40000 may be configured to provide such devices with rule sets or executable instructions to establish a connection to the WAN in response to or in expectation of an occurrence of the network-constrained state. For example, in response to network congestion or bandwidth exhaustion, one or more of the WAN-attachable devices may be actuated to establish additional connections to the digital twin system 40000 in parallel to the congested or exhausted connection (e.g., to provide additional communications bandwidth for connected devices).

[2117] Further, a reduction in data available for communication may inhibit the use of certain operational parameters. For example, a process may require suboptimal processing with lower data communication to prevent, for example, runaway of a reaction. The digital twin system 40000 may determine optimal parameters for a plurality of processes running at suboptimal levels by minimizing the time period to return to steady state after the network connectivity state has ceased.

[2118] Additionally or alternatively, the digital twin system 40000 may provide alerts to one or more users or system elements in response to detecting states, such as a network-constrained state. For example, actions taken by the digital twin system 40000 may be able to prevent any noticeable impact on the environment from a network-constrained state, so the digital twin system 40000 may

provide an alert to users of the digital twin system 40000. The alert may be a notification of an occurrence of the network-constrained state, an indication of data corresponding to the network-constrained state (e.g., reliability statistics or constrained points), instructions for reducing impact of future events (e.g., locations to increase connection points or available bandwidth), instructions on effect of the particular network-constrained state on the environment (e.g., missing data from affected devices), combinations thereof, and the like.

[2119]

[2120] In example embodiments, the digital twin I/O system 40004 includes a pathing module 400310. The pathing module 400310 may ingest navigational data from the elements 40302, provide and/or request navigational data to components of the digital twin system 40000 (e.g., the digital twin simulation system 40006, the digital twin behavior system 108, and/or the cognitive intelligence system 40010), and/or output navigational data to elements 40302 (e.g., to the wearable devices 40054). The navigational data may be collected or estimated using, for example, historical data, guidance data provided to the elements 40302, combinations thereof, and the like.

[2121] For example, the navigational data may be collected or estimated using historical data stored by the digital twin system 40000. The historical data may include or be processed to provide information such as acquisition time, associated elements 40302, polling intervals, task performed, laden or unladen conditions, whether prior guidance data was provided and/or followed, conditions of the environment 40020, other elements 40302 within the environment 40020, combinations thereof, and the like. The estimated data may be determined using one or more suitable pathing algorithms. For example, the estimated data may be calculated using suitable order-picking algorithms, suitable path-search algorithms, combinations thereof, and the like. The order-picking algorithm may be, for example, a largest gap algorithm, an s-shape algorithm, an aisle-by-aisle algorithm, a combined algorithm, combinations thereof, and the like. The path-search algorithms may be, for example, Dijkstra's algorithm, the A* algorithm, hierarchical path-finding algorithms, incremental path-finding algorithms, any angle path-finding algorithms, flow field algorithms, combinations thereof, and the like.

[2122] Additionally or alternatively, the navigational data may be collected or estimated using guidance data of the worker. The guidance data may include, for example, a calculated route provided to a device of the worker (e.g., a mobile device or the wearable device 40054). In another example, the guidance data may include a calculated route provided to a device of the worker that instructs the worker to collect vibration measurements from one or more locations on one or more machines along the route. The collected and/or estimated navigational data may be provided to a user of the digital twin system 40000 for visualization, used by other components of the digital twin system 40000 for analysis, optimization, and/or alteration, provided to one or more elements 40302, combinations thereof, and the like.

[2123] In example embodiments, the digital twin system 40000 ingests navigational data for a set of workers for representation in a digital twin. Additionally or alternatively, the digital twin system 40000 ingests navigational data for a set of mobile equipment assets of an industrial environment into a digital twin.

[2124] In example embodiments, the digital twin system 40000 ingests a system for modeling traffic of mobile elements in an industrial digital twin. For example, the digital twin system 40000 may model traffic patterns for workers or persons within the environment 40020, mobile equipment assets, combinations thereof, and the like. The traffic patterns may be estimated based on modeling traffic patterns from historical data and contemporaneous ingested data. Further, the traffic patterns may be continuously or intermittently updated depending on conditions within the environment 40020 (e.g., a plurality of autonomous mobile equipment assets may provide information to the digital twin system 40000 at a slower update interval than the environment 40020 including both workers and mobile equipment assets).

[2125] The digital twin system 40000 may alter traffic patterns (e.g., by providing updated navigational data to one or more of the mobile elements) to achieve one or more predetermined criteria. The predetermined criteria may include, for example, increasing process efficiency, decreasing interactions between laden workers and mobile equipment assets, minimizing worker path length, routing mobile equipment around paths or potential paths of persons, combinations thereof, and the like.

[2126] In example embodiments, the digital twin system 40000 may provide traffic data and/or navigational information to mobile elements in an industrial digital twin. The navigational information may be provided as instructions or rule sets, displayed path data, or selective actuation of devices. For example, the digital twin system 40000 may provide a set of instructions to a robot to direct the robot to and/or along a desired route for collecting vibration data from one or more specified locations on one or more specified machines along the route using a vibration sensor. The robot may communicate updates to the system including obstructions, reroutes, unexpected interactions with other assets within the environment 40020, etc.

[2127] In example embodiments, an ant-based system 40074 enables industrial entities, including robots, to lay a trail with one or more messages for other industrial entities, including themselves, to follow in later journeys. In example embodiments, the messages include information related to vibration measurement collection. In example embodiments, the messages include information related to vibration sensor measurement locations. In example embodiments, the trails may be configured to fade over time. In example embodiments, the ant-based trails may be experienced via an augmented reality system. In example embodiments, the ant-based trails may be experienced via a virtual reality system. In example embodiments, the ant-based trails may be experienced via a mixed reality system. In example embodiments, ant-based tagging of areas can trigger a pain-response and/or accumulate into a warning signal. In example embodiments, the ant-based trails may be configured to generate an information filtering response. In example embodiments, the ant-based trails may be configured to generate an information filtering response wherein the information filtering response is a heightened sense of visual awareness. In example embodiments, the ant-based trails may be configured to generate an information filtering response wherein the information filtering response is a heightened sense of acoustic awareness. In example embodiments, the messages include vectorized data.

[2128] In example embodiments, the digital twin system 40000 includes design specification information for representing a real-world element 40302r using a digital twin 40302d. The digital may correspond to an existing real-world element 40302r or a potential real-world element 40302r. The design specification information may be received from one or more sources. For example, the design specification information may include design parameters set by user input, determined by the digital twin system 40000 (e.g., the via digital twin simulation system 40006), optimized by users or the digital twin simulation system 40006, combinations thereof, and the like. The digital twin simulation system 40006 may represent the design specification information for the component to users, for example, via a display device or a wearable device. The design specification information may be displayed schematically (e.g., as part of a process diagram or table of information) or as part of an augmented reality or virtual reality display. The design specification information may be displayed, for example, in response to a user interaction with the digital twin system 40000 (e.g., via user selection of the element or user selection to generally include design specification information within displays). Additionally or alternatively, the design specification information may be displayed automatically, for example, upon the element coming within view of an augmented reality or virtual reality device. In example embodiments, the displayed design specification information may further include indicia of information source (e.g., different displayed colors indicate user input versus digital twin system 40000 determination), indicia of mismatches (e.g., between design specification information and operational information), combinations thereof, and the like.

[2129] In example embodiments, the digital twin system 40000 embeds a set of control instructions for a wearable device within an industrial digital twin, such that the control instructions are provided to the wearable device to induce an experience for a wearer of the wearable device upon interaction with an element of the industrial digital twin. The induced experience may be, for example, an augmented reality experience or a virtual reality experience. The wearable device, such as a headset, may be configured to output video, audio, and/or haptic feedback to the wearer to induce the experience. For example, the wearable device may include a display device and the experience may include display of information related to the respective digital twin. The information displayed may include maintenance data associated with the digital twin, vibration data associated with the digital twin, vibration measurement location data associated with the digital twin, financial data associated with the digital twin, such as a profit or loss associated with operation of the digital twin, manufacturing KPIs associated with the digital twin, information related to an occluded element (e.g., a sub-assembly) that is at least partially occluded by a foreground element (e.g., a housing), a virtual model of the occluded element overlaid on the occluded element and visible with the foreground element, operating parameters for the occluded element, a comparison to a design parameter corresponding to the operating parameter displayed, combinations thereof, and the like. Comparisons may include, for example, altering display of the operating parameter to change a color, size, and/or display period for the operating parameter.

[2130] In example embodiments, the displayed information may include indicia for removable elements that are or may be configured to provide access to the occluded element with each

indicium being displayed proximate to or overlying the respective removable element. Further, the indicia may be sequentially displayed such that a first indicium corresponding to a first removable element (e.g., a housing) is displayed, and a second indicium corresponding to a second removable element (e.g., an access panel within the housing) is displayed in response to the worker removing the first removable element. In example embodiments, the induced experience allows the wearer to see one or more locations on a machine for optimal vibration measurement collection. In an example, the digital twin system 40000 may provide an augmented reality view that includes highlighted vibration measurement collection locations on a machine and/or instructions related to collecting vibration measurements. Furthering the example, the digital twin system 40000 may provide an augmented reality view that includes instructions related to timing of vibration measurement collection. Information utilized in displaying the highlighted placement locations may be obtained using information stored by the digital twin system 40000. In example embodiments, mobile elements may be tracked by the digital twin system 40000 (e.g., via observational elements within the environment 40020 and/or via pathing information communicated to the digital twin system 40000) and continually displayed by the wearable device within the occluded view of the worker. This optimizes movement of elements within the environment 40020, increases worker safety, and minimizes down time of elements resulting from damage.

[2131] In example embodiments, the digital twin system 40000 may provide an augmented reality view that displays mismatches between design parameters or expected parameters of real-world elements 40302r to the wearer. The displayed information may correspond to real-world elements 40302r that are not within the view of the wearer (e.g., elements within another room or obscured by machinery). This allows the worker to quickly and accurately troubleshoot mismatches to determine one or more sources for the mismatch. The cause of the mismatch may then be determined, for example, by the digital twin system 40000 and corrective actions ordered. In example embodiments, a wearer may be able to view malfunctioning subcomponents of machines without removing occluding elements (e.g., housings or shields). Additionally or alternatively, the wearer may be provided with instructions to repair the device, for example, including display of the removal process (e.g., location of fasteners to be removed), assemblies or subassemblies that should be transported to other areas for repair (e.g., dust-sensitive components), assemblies or subassemblies that need lubrication, and locations of objects for reassembly (e.g., storing location that the wearer has placed removed objects and directing the wearer or another wearer to the stored locations to expedite reassembly and minimize further disassembly or missing parts in the reassembled element). This can expedite repair work, minimize process impact, allow workers to disassemble and reassemble equipment (e.g., by coordinating disassembly without direct communication between the workers), increase equipment longevity and reliability (e.g., by assuring that all components are properly replaced prior to placing back in service), combinations thereof, and the like.

[2132] In example embodiments, the induced experience includes a virtual reality view or an augmented reality view that allows the wearer to see information related to existing or planned

elements. The information may be unrelated to physical performance of the element (e.g., financial performance such as asset value, energy cost, input material cost, output material value, legal compliance, and corporate operations). One or more wearers may perform a virtual walkthrough or an augmented walkthrough of the industrial environment 40020.

[2133] In examples, the wearable device displays compliance information that expedites inspections or performance of work. For example, an electrical inspector may walk through a site and inspect obscured connections for compliance with particular codes even when objects obscure the relevant inspection points (e.g., equipment or finish materials). This expedites construction and inspection and minimizes change orders because further work does not need to be delayed or altered to wait for inspector approval of the completed work. Further, this minimizes rework as compliance may be ensured by persons unfamiliar with the code (e.g., a worker unfamiliar with electrical code may be able to ensure compliance of the electrical work prior to placement of equipment).

[2134] In further examples, the wearable device displays financial information that is used to identify targets for alteration or optimization. For example, a manager or officer may inspect the environment 40020 for compliance with updated regulations, including information regarding compliance with former regulations, "grandfathered," and/or excepted elements. This can be used to reduce unnecessary downtime (e.g., scheduling upgrades for the least impactful times, such as during planned maintenance cycles), prevent unnecessary upgrades (e.g., replacing grandfathered or excepted equipment), and reduce capital investment.

[2135] Referring back to FIG. 213, in example embodiments, the digital twin system 40000 may include, integrate, integrate with, manage, handle, link to, take input from, provide output to, control, coordinate with, or otherwise interact with a digital twin dynamic model system 40008. The digital twin dynamic model system 40008 can update the properties of a set of digital twins of a set of industrial entities and/or environments, including properties of physical industrial assets, workers, processes, manufacturing facilities, warehouses, and the like (or any of the other types of entities or environments described in this disclosure or in the documents incorporated by reference herein) in such a manner that the digital twins may represent those industrial entities and environments, and properties or attributes thereof, in real-time or very near real-time. In example embodiments, the digital twin dynamic model system 40008 may obtain sensor data received from a sensor system 40030 and may determine one or more properties of an industrial environment or an industrial entity within an environment based on the sensor data and based on one or more dynamic models.

[2136] In example embodiments, the digital twin dynamic model system 40008 may update/assign values of various properties in a digital twin and/or one or more embedded digital twins, including, but not limited to, vibration values, vibration fault level states, probability of failure values, probability of downtime values, cost of downtime values, probability of shutdown values, financial values, KPI values, temperature values, humidity values, heat flow values, fluid flow values, radiation values, substance concentration values, velocity values, acceleration values, location

values, pressure values, stress values, strain values, light intensity values, sound level values, volume values, shape characteristics, material characteristics, and dimensions.

[2137] In example embodiments, a digital twin may be comprised of (e.g., via reference) of other embedded digital twins. For example, a digital twin of a manufacturing facility may include an embedded digital twin of a machine and one or more embedded digital twins of one or more respective motors enclosed within the machine. A digital twin may be embedded, for example, in the memory of an industrial machine that has an onboard IT system (e.g., the memory of an Onboard Diagnostic System, control system (e.g., SCADA system) or the like). Other non-limiting examples of where a digital twin may be embedded include the following: on a wearable device of a worker; in memory on a local network asset, such as a switch, router, access point, or the like; in a cloud computing resource that is provisioned for an environment or entity; and on an asset tag or other memory structure that is dedicated to an entity.

[2138] In one example, the digital twin dynamic model system 40008 can update vibration characteristics throughout an industrial environment digital twin based on captured vibration sensor data measured at one or more locations in the industrial environment and one or more dynamic models that model vibration within the industrial environment digital twin. The industrial digital twin may, before updating, already contain information about properties of the industrial entities and/or environment that can be used to feed a dynamic model, such as information about materials, shapes/volumes (e.g., of conduits), positions, connections/interfaces, and the like, such that vibration characteristics can be represented for the entities and/or environment in the digital twin. Alternatively, the dynamic models may be configured using such information. Alternatively, the thermodynamic models may be configured using such information. Other sensor data may also work to update thermodynamic behavior, such as pressure data (e.g., using PV=nRT). Thermodynamic models may also be configured to represent the diffusion of heat through static objects (e.g., big metal machines) as well as through fluids (e.g., circulating fluids in a cooling system).

[2139] In another example, the digital twin dynamic system 40008 can update the concentration values for a chemical compound (analyte) throughout an industrial environment digital twin based on captured chemical sensor data and one or more diffusion models that model the concentrations of chemicals within the industrial environment digital twin. The industrial environment digital twin can include a set of properties and/or attributes of the environment and/or entities that can help supply inputs to a chemical diffusion model and/or chemical interaction/reaction model, such as chemical compositions of materials, fluids, gases, etc., shapes/volumes of components, conduits, spaces, etc., temperatures and pressures, and other factors. The sensors can be chemical sensors, but also pressure, temperature, flow and other sensors that may inform the diffusion model.

[2140] In example embodiments, the digital twin dynamic model system 40008 can update the properties of a digital twin and/or one or more embedded digital twins on behalf of a client application 40070. In example embodiments, a client application 40070 may be an application relating to an industrial component or environment (e.g., monitoring an industrial facility or a component within, simulating an industrial environment, or the like). In example embodiments,

the client application 40070 may be used in connection with both fixed and mobile data collection systems. In example embodiments, the client application 40070 may be used in connection with Industrial Internet of Things sensor system 40030.

[2141] In example embodiments, the digital twin dynamic model system 40008 leverages digital twin dynamic models 400100 to model the behavior of an industrial entity and/or environment. Dynamic models 400100 may enable digital twins to represent physical reality, including the interactions of industrial entities, by using a limited number of measurements to enrich the digital representation of an industrial entity and/or environment, such as based on scientific principles. In example embodiments, the dynamic models 400100 are formulaic or mathematical models. In example embodiments, the dynamic models 400100 adhere to scientific laws, laws of nature, and formulas (e.g., Newton's laws of motion, second law of thermodynamics, Bernoulli's principle, ideal gas law, Dalton's law of partial pressures, Hooke's law of elasticity, Fourier's law of heat conduction, Archimedes' principle of buoyancy, and the like). In example embodiments, the dynamic models are machine-learned models. For example, temperature sensors in a warehouse may each take a temperature measurement at specific geospatial coordinates, but these limited measurements do not give values for the other locations in the warehouse, such as where there is no sensor coverage. In this example, the dynamic models can be configured to model temperatures throughout the warehouse using the limited number of sensor measurements to provide a more enriched representation of the warehouse digital twin.

[2142] In example embodiments, the digital twin system 40000 may have a digital twin dynamic model datastore 400102 for storing dynamic models 400100 that may be represented in digital twins. In example embodiments, digital twin dynamic model datastore can be searchable and/or discoverable. In example embodiments, digital twin dynamic model datastore can contain metadata that allows a user to understand what characteristics a given dynamic model can handle, what inputs are required, what outputs are provided, and the like. In example embodiments, digital twin dynamic model datastore 400102 can be hierarchical (such as where a model can be deepened or made more simple based on the extent of available data and/or inputs, the granularity of the inputs, and/or situational factors (such as where something becomes of high interest and a higher fidelity model is accessed for a period of time).

[2143] In example embodiments, a digital twin or digital representation of an industrial entity or facility may include a set of data structures that collectively define a set of properties of a represented physical industrial asset, device, worker, process, facility, and/or environment, and/or possible behaviors thereof. In example embodiments, the digital twin dynamic model system 40008 may leverage the dynamic models 400100 to inform the set of data structures that collectively define a digital twin with real-time data values. The digital twin dynamic models 400100 may receive one or more sensor measurements, Industrial Internet of Things device data, and/or other suitable data as inputs and calculate one or more outputs based on the received data and one or more dynamic models 400100. The digital twin dynamic model system 40008 then uses the one or more outputs to update the digital twin data structures.

[2144] In one example, the set of properties of a digital twin of an industrial asset that may be updated by the digital twin dynamic model system 40008 using dynamic models 400100 may include the vibration characteristics of the asset, temperature(s) of the asset, the state of the asset (e.g., a solid, liquid, or gas), the location of the asset, the displacement of the asset, the velocity of the asset, the acceleration of the asset, probability of downtime values associated with the asset, cost of downtime values associated with the asset, probability of shutdown values associated with the asset, manufacturing KPIs associated with the asset, financial information associated with the asset, heat flow characteristics associated with the asset, fluid flow rates associated with the asset (e.g., fluid flow rates of a fluid flowing through a pipe), identifiers of other digital twins embedded within the digital twin of the asset and/or identifiers of digital twins embedding the digital twin of the asset, and/or other suitable properties. Dynamic models 400100 associated with a digital twin of an asset can be configured to calculate, interpolate, extrapolate, and/or output values for such asset digital twin properties based on input data collected from sensors and/or devices disposed in the industrial setting and/or other suitable data and subsequently populate the asset digital twin with the calculated values.

[2145] In example embodiments, the set of properties of a digital twin of an industrial device that may be updated by the digital twin dynamic model system 40008 using dynamic models 400100 may include the status of the device, a location of the device, the temperature(s) of a device, a trajectory of the device, identifiers of other digital twins that the digital twin of the device is embedded within, embeds, is linked to, includes, integrates with, takes input from, provides output to, and/or interacts with and the like. Dynamic models 400100 associated with a digital twin of a device can be configured to calculate or output values for these device digital twin properties based on input data and subsequently update the device digital twin with the calculated values.

[2146] In example embodiments, the set of properties of a digital twin of an industrial worker that may be updated by the digital twin dynamic model system 40008 using dynamic models 400100 may include the status of the worker, the location of the worker, a stress measure for the worker, a task being performed by the worker, a performance measure for the worker, and the like. Dynamic models associated with a digital twin of an industrial worker can be configured to calculate or output values for such properties based on input data, which then may be used to populate industrial worker digital twin. In example embodiments, industrial worker dynamic models (e.g., psychometric models) can be configured to predict reactions to stimuli, such as cues that are given to workers to direct them to undertake tasks and/or alerts or warnings that are intended to induce safe behavior. In example embodiments, industrial worker dynamic models may be workflow models (Gantt charts and the like), failure mode effects analysis models (FMEA), biophysical models (such as to model worker fatigue, energy utilization, hunger), and the like.

[2147] Example properties of a digital twin of an industrial environment that may be updated by the digital twin dynamic model system 40008 using dynamic models 400100 may include the dimensions of the industrial environment, the temperature(s) of the industrial environment, the humidity value(s) of the industrial environment, the fluid flow characteristics in the industrial environment, the heat flow characteristics of the industrial environment, the lighting characteristics

of the industrial environment, the acoustic characteristics of the industrial environment the physical objects in the environment, processes occurring in the industrial environment, currents of the industrial environment (if a body of water), and the like. Dynamic models associated with a digital twin of an industrial environment can be configured to calculate or output these properties based on input data collected from sensors and/or devices disposed in the industrial environment and/or other suitable data and subsequently populate the industrial environment digital twin with the calculated values.

[2148] In example embodiments, dynamic models 400100 may adhere to physical limitations that define boundary conditions, constants or variables for digital twin modeling. For example, the physical characterization of a digital twin of an industrial entity or industrial environment may include a gravity constant (e.g., 9.8 m/s^2), friction coefficients of surfaces, thermal coefficients of materials, maximum temperatures of assets, maximum flow capacities, and the like. Additionally or alternatively, the dynamic models may adhere to the laws of nature. For example, dynamic models may adhere to the laws of thermodynamics, laws of motion, laws of fluid dynamics, laws of buoyancy, laws of heat transfer, laws or radiation, laws of quantum dynamics, and the like. In example embodiments, dynamic models may adhere to biological aging theories or mechanical aging principles. Thus, when the digital twin dynamic model system 40008 facilitates a real-time digital representation, the digital representation may conform to dynamic models, such that the digital representations mimic real world conditions. In example embodiments, the output(s) from a dynamic model can be presented to a human user and/or compared against real-world data to ensure convergence of the dynamic models with the real world. Furthermore, as dynamic models are based partly on assumptions, the properties of a digital twin may be improved and/or corrected when a real-world behavior differs from that of the digital twin. In example embodiments, additional data collection and/or instrumentation can be recommended based on the recognition that an input is missing from a desired dynamic model, that a model in operation isn't working as expected (perhaps due to missing and/or faulty sensor information), that a different result is needed (such as due to situational factors that make something of high interest), and the like.

[2149] Dynamic models may be obtained from a number of different sources. In example embodiments, a user can upload a model created by the user or a third party. Additionally or alternatively, the models may be created on the digital twin system using a graphical user interface. The dynamic models may include bespoke models that are configured for a particular environment and/or set of industrial entities and/or agnostic models that are applicable to similar types of digital twins. The dynamic models may be machine-learned models.

[2150] FIG. 217 illustrates example embodiments of a method at 41100 for updating a set of properties of a digital twin and/or one or more embedded digital twins on behalf of client applications 40070. In example embodiments, digital twin dynamic model system 40008 leverages one or more dynamic models 400100 to update a set of properties of a digital twin and/or one or more embedded digital twins on behalf of client application 40070 based on the impact of collected sensor data from sensor system 40030, data collected from Internet of Things connected devices 40024, and/or other suitable data in the set of dynamic models 400100 that are used to enable the

industrial digital twins. In example embodiments, the digital twin dynamic model system 40008 may be instructed to run specific dynamic models using one or more digital twins that represent physical industrial assets, devices, workers, processes, and/or industrial environments that are managed, maintained, and/or monitored by the client applications 40070.

[2151] In example embodiments, the digital twin dynamic model system 40008 may obtain data from other types of external data sources that are not necessarily industrial-related data sources, but may provide data that can be used as input data for the dynamic models. For example, weather data, news events, social media data, and the like may be collected, crawled, subscribed to, and the like to supplement sensor data, Industrial Internet of Things device data, and/or other data that is used by the dynamic models. In example embodiments, the digital twin dynamic model system 40008 may obtain data from a machine vision system. Machine vision system, which may be included in the sensor system 40030 and the video sensors 40052, may use video and/or still images to provide measurements (e.g., locations, statuses, and the like) that may be used as inputs by the dynamic models.

[2152] In example embodiments, the digital twin dynamic model system 40008 may feed this data into one or more of the dynamic models discussed above to obtain one or more outputs. These outputs may include calculated vibration fault level states, vibration severity unit values, vibration characteristics, probability of failure values, probability of downtime values, probability of shutdown values, cost of downtime values, cost of shutdown values, time to failure values, temperature values, pressure values, humidity values, precipitation values, visibility values, air quality values, strain values, stress values, displacement values, velocity values, acceleration values, location values, performance values, financial values, manufacturing KPI values, electrodynamic values, thermodynamic values, fluid flow rate values, and the like. The client application 40070 may then initiate a digital twin visualization event using the results obtained by the digital twin dynamic model system 40008. In example embodiments, the visualization may be a heat map visualization.

[2153] In example embodiments, the digital twin dynamic model system 40008 may receive requests to update one or more properties of digital twins of industrial entities and/or environments such that the digital twins represent the industrial entities and/or environments in real-time. At 41102, the digital twin dynamic model system 40008 receives a request to update one or more properties of one or more of the digital twins of industrial entities and/or environments. For example, the digital twin dynamic model system 40008 may receive the request from a client application 40070 or from another process executed by the digital twin system 40000 (e.g., a predictive maintenance process). The request may indicate the one or more properties and the digital twin or digital twins implicated by the request. At 41104, the digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins, including any embedded digital twins, from digital twin datastore 40016. At 41108, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from digital twin dynamic model store. At 41110, the digital twin dynamic model system

40008 selects one or more sensors from sensor system 40030, data collected from Internet of Things connected devices 40024, and/or other data sources from digital twin I/O system 40004 based on available data sources and the one or more required inputs of the dynamic model(s). In example embodiments, the data sources may be defined in the inputs required by the one or more dynamic models or may be selected using a lookup table. At 41112, the digital twin dynamic model system 40008 retrieves the selected data from digital twin I/O system 40004. At 41114, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved input data (e.g., vibration sensor data, Industrial Internet of Things device data, and the like) as inputs and determines one or more output values based on the dynamic model(s) and the input data. At 41120, the digital twin dynamic model system 40008 updates the values of one or more properties of the one or more digital twins based on the one or more outputs of the dynamic model(s).

[2154] In example embodiments, client application 40070 may be configured to provide a digital representation and/or visualization of the digital twin of an industrial entity. In example embodiments, the client application 40070 may include one or more software modules that are executed by one or more server devices. These software modules may be configured to quantify properties of the digital twin, model properties of a digital twin, and/or to visualize digital twin behaviors. In example embodiments, these software modules may enable a user to select a particular digital twin behavior visualization for viewing. In example embodiments, these software modules may enable a user to select to view a digital twin behavior visualization playback. In example embodiments, the client application 40070 may provide a selected behavior visualization to digital twin dynamic model system 40008.

[2155] In example embodiments, the digital twin dynamic model system 40008 may receive requests from client application 40070 to update properties of a digital twin in order to enable a digital representation of an industrial entity and/or environment wherein the real-time digital representation is a visualization of the digital twin. In example embodiments, a digital twin may be rendered by a computing device, such that a human user can view the digital representations of real-world industrial assets, devices, workers, processes and/or environments. For example, the digital twin may be rendered and outcome to a display device. In example embodiments, dynamic model outputs and/or related data may be overlaid on the rendering of the digital twin. In example embodiments, dynamic model outputs and/or related information may appear with the rendering of the digital twin in a display interface. In example embodiments, the related information may include real-time video footage associated with the real-world entity represented by the digital twin. In example embodiments, the related information may include a sum of each of the vibration fault level states in the machine. In example embodiments, the related information may be graphical information. In example embodiments, the graphical information may depict motion and/or motion as a function of frequency for individual machine components. In example embodiments, graphical information may depict motion and/or motion as a function of frequency for individual machine components, wherein a user is enabled to select a view of the graphical information in the x, y, and z dimensions. In example embodiments, graphical information may depict motion and/or motion as a function of frequency for individual machine components,

wherein the graphical information includes harmonic peaks and peaks. In example embodiments, the related information may be cost data, including the cost of downtime per day data, cost of repair data, cost of new part data, cost of new machine data, and the like. In example embodiments, related information may be a probability of downtime data, probability of failure data, and the like. In example embodiments, related information may be time to failure data.

[2156] In example embodiments, the related information may be recommendations and/or insights. For example, recommendations or insights received from the cognitive intelligence system related to a machine may appear with the rendering of the digital twin of a machine in a display interface.

[2157] In example embodiments, clicking, touching, or otherwise interacting with the digital twin rendered in the display interface can allow a user to "drill down" and see underlying subsystems or processes and/or embedded digital twins. For example, in response to a user clicking on a machine bearing rendered in the digital twin of a machine, the display interface can allow a user to drill down and see information related to the bearing, view a 3D visualization of the bearing's vibration, and/or view a digital twin of the bearing.

[2158] In example embodiments, clicking, touching, or otherwise interacting with information related to the digital twin rendered in the display interface can allow a user to "drill down" and see underlying information.

[2159] FIG. 218 illustrates example embodiments of a display interface that renders the digital twin of a dryer centrifuge and other information related to the dryer centrifuge.

[2160] In example embodiments, the digital twin may be rendered and output in a virtual reality display. For example, a user may view a 3D rendering of an environment (e.g., using a monitor or a virtual reality headset). The user may also inspect and/or interact with digital twins of industrial entities. In example embodiments, a user may view processes being performed with respect to one or more digital twins (e.g., collecting measurements, movements, interactions, inventorying, loading, packing, shipping, and the like). In example embodiments, a user may provide input that controls one or more properties of a digital twin via a graphical user interface.

[2161] In example embodiments, the digital twin dynamic model system 40008 may receive requests from client application 40070 to update properties of a digital twin in order to enable a digital representation of industrial entities and/or environments wherein the digital representation is a heat map visualization of the digital twin. In example embodiments, a platform is provided having heat maps displaying collected data from the sensor system 40030, Internet of Things connected devices 40024, and data outputs from dynamic models 400100 for providing input to a display interface. In example embodiments, the heat map interface is provided as an output for digital twin data, such as for handling and providing information for visualization of various sensor data, dynamic model output data, and other data (such as map data, analog sensor data, and other data), such as to another system, such as a mobile device, tablet, dashboard, computer, AR/VR device, or the like. A digital twin representation may be provided in a form factor (e.g., user device, VR-enabled device, AR-enabled device, or the like) suitable for delivering visual input to a user, such as the presentation of a map that includes indicators of levels of analog sensor data, digital

sensor data, and output values from the dynamic models (such as data indicating vibration fault level states, vibration severity unit values, probability of downtime values, cost of downtime values, probability of shutdown values, time to failure values, probability of failure values, manufacturing KPIs, temperatures, levels of rotation, vibration characteristics, fluid flow, heating or cooling, pressure, substance concentrations, and many other output values). In example embodiments, signals from various sensors or input sources (or selective combinations, permutations, mixes, and the like) as well as data determined by the digital twin dynamic model system 40008 may provide input data to a heat map. Coordinates may include real world location coordinates (such as geo-location or location on a map of an environment), as well as other coordinates, such as time-based coordinates, frequency-based coordinates, or other coordinates that allow for representation of analog sensor signals, digital signals, dynamic model outputs, input source information, and various combinations, in a map-based visualization, such that colors may represent varying levels of input along the relevant dimensions. For example, among many other possibilities, if an industrial machine component is at a critical vibration fault level state, the heat map interface may alert a user by showing the machine component in orange. In the example of a heat map, clicking, touching, or otherwise interacting with the heat map can allow a user to drill down and see underlying sensor, dynamic model outputs, or other input data that is used as an input to the heat map display. In other examples, such as ones where a digital twin is displayed in a VR or AR environment, if an industrial machine component is vibrating outside of normal operation (e.g., at a suboptimal, critical, or alarm vibration fault level), a haptic interface may induce vibration when a user touches a representation of the machine component, or if a machine component is operating in an unsafe manner, a directional sound signal may direct a user's attention toward the machine in digital twin, such as by playing in a particular speaker of a headset or other sound system.

[2162] In example embodiments, the digital twin dynamic model system 40008 may take a set of ambient environmental data and/or other data and automatically update a set of properties of a digital twin of an industrial entity or facility based on the impact of the environmental data and/or other data in the set of dynamic models 400100 that are used to enable the digital twin. Ambient environmental data may include temperature data, pressure data, humidity data, wind data, rainfall data, tide data, storm surge data, cloud cover data, snowfall data, visibility data, water level data, and the like. Additionally or alternatively, the digital twin dynamic model system 40008 may use a set of environmental data measurements collected by a set of Internet of Things connected devices 40024 disposed in an industrial setting as inputs for the set of dynamic models 400100 that are used to enable the digital twin. For example, digital twin dynamic model system 40008 may feed the dynamic models 400100 data collected, handled or exchanged by Internet of Things connected devices 40024, such as cameras, monitors, embedded sensors, mobile devices, diagnostic devices and systems, instrumentation systems, telematics systems, and the like, such as for monitoring various parameters and features of machines, devices, components, parts, operations, functions, conditions, states, events, workflows and other elements (collectively encompassed by the term "states") of industrial environments. Other examples of Internet of

Things connected devices include smart fire alarms, smart security systems, smart air quality monitors, smart/learning thermostats, and smart lighting systems.

[2163] FIG. 219 illustrates example embodiments of a method at 42000 for updating a set of vibration fault level states for a set of bearings in a digital twin of a machine. In this example, a client application 40070, which interfaces with digital twin dynamic model system 40008, may be configured to provide a visualization of the fault level states of the bearings in the digital twin of the machine.

[2164] In this example, the digital twin dynamic model system 40008 may receive requests from client application 40070 to update the vibration fault level states of the machine digital twin. At 42002, digital twin dynamic model system 40008 receives a request from client application 40070 to update one or more vibration fault level states of the machine digital twin. Next, at 42004, digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore 40016. In this example, the digital twin dynamic model system 40008 may retrieve the digital twin of the machine and any embedded digital twins, such as any embedded motor digital twins and bearing digital twins, and any digital twins that embed the machine digital twin, such as the manufacturing facility digital twin. At 42008, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from the digital twin dynamic model datastore 400102. At 42010, the digital twin dynamic model system 40008 selects dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) via digital twin I/O system 40004 based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the and the one or more required inputs of the dynamic model(s). In the present example, the retrieved dynamic model(s) 400100 may take one or more vibration sensor measurements from vibration sensors 40036 as inputs to the dynamic models. In example embodiments, vibration sensors 40036 may be optical vibration sensors, single axis vibration sensors, tri-axial vibration sensors, and the like. At 42012, digital twin dynamic model system 40008 retrieves one or more measurements from each of the selected data sources from the digital twin I/O system 40004. Next, at 42014, digital twin dynamic model system 40008 runs the dynamic model(s), using the retrieved vibration sensor measurements as inputs, and calculates one or more outputs that represent bearing vibration fault level states. Next, at 42018, the digital twin dynamic model system 40008 updates one or more bearing fault level states of the manufacturing facility digital twin, machine digital twin, motor digital twin, and/or bearing digital twins based on the one or more outputs of the dynamic model(s). The client application 40070 may obtain vibration fault level states of the bearings and may display the obtained vibration fault level state associated with each bearing and/or display colors associated with fault level severity (e.g., red for alarm, orange for critical, yellow for suboptimal, green for normal operation) in the rendering of one or more of the digital twins on a display interface.

[2165] Taking the example further, additionally or alternatively, client application 40070 may be configured to provide a heat map visualization of strain on industrial entities within the

manufacturing facility, such as a pipe. Piping materials can exhibit a linear expansion and contraction with temperature and thermal pipe expansion may cause strain on piping materials.

[2166] The rate of thermal expansion and contraction is characterized by the coefficient of thermal expansion. The change in dimensions of the pipe could be defined by:

[2167] $\varepsilon = a\,(T2-T1)$ (Equation 1)

[2168] where:

[2169] ε = strain (in/in)

[2170] a = Coefficient of thermal expansion (in/in-°F)

[2171] T2 = End temperature (°F)

[2172] T1 = Starting temperature (°F)

[2173] Given the temperature at installation (T1), coefficient of thermal expansion, and a sensor measurement giving the real-time temperature for a particular point on a pipe (T2), the pipe strain values may be calculated from dynamic models that take one or more temperature measurements from temperature sensors 40032 as input(s) to the dynamic models and calculate one or more estimated pipe strain values in adherence to Equation 1. Additionally or alternatively, the dynamic models may be configured to take other suitable data as inputs (e.g., humidity data from humidity sensor 40034, pressure data from pressure sensor 40046, data from Internet of Things connected devices 40024, and the like) to calculate one or more pipe strain values. The digital twin dynamic system 40008 may then update the manufacturing facility digital twin, pipe digital twin, and any other suitable industrial entity digital twins with pipe strain values.

[2174] In another example, a client application 40070 may be an augmented reality application. In some embodiments of this example, the client application 40070 may obtain vibration fault level states of bearings in a field of view of an AR-enabled device (e.g., smart glasses) hosting the client application from the digital twin of the industrial environment (e.g., via an API, webhook, etc. of the digital twin system 40000) and may display the obtained vibration fault level states on the display of the AR-enabled device, such that the vibration fault level state displayed corresponds to the location in the field of view of the AR-enabled device. In this way, a vibration fault level state may be displayed even if there are no vibration sensors located within the field of view of the AR-enabled device.

[2175] FIG. 220 illustrates example embodiments of a method at 42100 for updating a set of vibration severity unit values of bearings in a digital twin of a machine. Vibration severity units may be measured as displacement, velocity, and acceleration.

[2176] In this example, client application 40070 that interfaces with the digital twin dynamic model system 40008 may be configured to provide a visualization of the three-dimensional vibration characteristics of bearings in a digital twin of a machine. RF spectrum characteristics may include signal frequency, signal amplitude, power level, and the like. In example embodiments, these characteristics may be measured with RF sensor 40078. RF sensors 40078 may be spectrum analyzers, a power meters, frequency counters, RF vector network analyzers (VNAs), and the like.

[2177] In this example, the digital twin dynamic model system 40008 may receive requests from client application 40070 to update the vibration severity unit values for bearings in the digital twin of a machine. At 42102, digital twin dynamic model system 40008 receives a request from client application 40070 to update one or more vibration severity unit value(s) of the manufacturing facility digital twin. Next, at 42104, digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore 40016. In this example, the digital twin dynamic model system 40008 may retrieve the digital twin of the machine and any embedded digital twins (e.g., digital twins of bearings and other components). At 42108, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore 400102. At 42110, the digital twin dynamic model system 40008 selects dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) via digital twin I/O system 40004 based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the one or more required inputs of the dynamic model(s). In the present example, the retrieved dynamic models may be configured to take one or more vibration sensor measurements as inputs and provide severity unit values for bearings in the machine. At 42112, digital twin dynamic model system 40008 retrieves one or more measurements from each of the selected sensors. In the present example, the digital twin dynamic model system 40008 retrieves measurements from vibration sensors 40036 via digital twin I/O system 40004. At 42114, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved vibration measurements as inputs and calculates one or more output values that represent vibration severity unit values for bearings in the machine. Next, at 42118, the digital twin dynamic model system 40008 updates one or more vibration severity unit values of the bearings in the machine digital twin and all other embedded digital twins or digital twins that embed the machine digital twin based on the one or more values output by the dynamic model(s).

[2178] FIG. 221 illustrates example embodiments of a method 42200 for updating a set of probability of failure values for machine components in the digital twin of a machine. FIG. 217 illustrates an example embodiment of a method for updating a set of electrodynamics-related values in the digital twin of an industrial environment such as a manufacturing facility. In this example, a client application 40070 that interfaces with the digital twin dynamic system 40008 may be configured to provide a visualization of the geospatial radiation characteristics of the manufacturing facility in the digital twin of the manufacturing facility. In example embodiments, the electrodynamics-related values may be related to electromagnetic field (EMF) radiation. Example types of EMF radiation include radio frequency, magnetic fields, and electrical fields. Geospatial radiation characteristics may include strength of radiation, frequency of radiation, and the like.

[2179] In this example, the digital twin dynamic model system 40008 may receive requests from client application 40070 to update the probability of failure values for components in a machine digital twin. At 42202, digital twin dynamic model system 40008 receives a request from client

application 40070 to update one or more probability of failure value(s) of the machine digital twin, any embedded component digital twins, and any digital twins that embed the machine digital twin such as a manufacturing facility digital twin. Next, at 42204, digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins. In this example, the digital twin dynamic model system 40008 may retrieve the digital twin of the manufacturing facility, the digital twin of the machine, and the digital twins of machine components from digital twin datastore 40016. At 42208, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore 400102. At 42210, the digital twin dynamic model system 40008 selects, via digital twin I/O system 40004, dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the and the one or more required inputs of the dynamic model(s). In the present example, the retrieved dynamic models may take one or more vibration measurements from vibration sensors 40036 and historical failure data as dynamic model inputs and output probability of failure values for the machine components in the digital twin of the machine. At 42212, digital twin dynamic model system 40008 retrieves data from each of the selected sensors and/or Internet of Things connected devices via digital twin I/O system 40004. At 42214, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved vibration data and historical failure data as inputs and calculates one or more outputs that represent probability of failure values for bearings in the machine digital twin. Next, at 42218, the digital twin dynamic model system 40008 updates one or more probability of failure values of the bearings in the machine digital twin, all embedded digital twins, and all digital twins that embed the machine digital twin based on the output of the dynamic model(s).

[2180] FIG. 222 illustrates example embodiments of a method 42300 for updating a set of probability of downtime for machines in the digital twin of a manufacturing facility. Chemical characteristics may include chemicals present in an environment, chemical concentrations, and the like. Chemical sensors 40054 may detect and measure the concentration of target molecules (also known as analytes). In example embodiments, chemical sensors 40054 may be gas sensors (such as semiconductor gas sensors, electrochemical gas sensors, contact combustion gas sensors, optical gas sensors, and polymer gas sensors), ion sensors, and humidity sensors.

[2181] In this example, client application 40070, which interfaces with the digital twin dynamic model system 40008, may be configured to provide a visualization of the probability of downtime values of a manufacturing facility in the digital twin of the manufacturing facility.

[2182] In this example, the digital twin dynamic model system 40008 may receive requests from client application 40070 to assign probability of downtime values to machines in a manufacturing facility digital twin. At 42302, digital twin dynamic model system 40008 receives a request from client application 40070 to update one or more probability of downtime values of machines in the manufacturing facility digital twin and any embedded digital twins such as the individual machine

digital twins. Next, at 42304, digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore 40016. In this example, the digital twin dynamic model system 40008 may retrieve the digital twin of the manufacturing facility and any embedded digital twins from digital twin datastore 40016. At 42308, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore 400102. At 42310, the digital twin dynamic model system 40008 selects dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system 40004. In the present example, the dynamic model(s) may be configured to take vibration measurements from vibration sensors and historical downtime data as inputs and output probability of downtime values for different machines throughout the manufacturing facility. At 42312, digital twin dynamic model system 40008 retrieves one or more measurements from each of the selected sensors via digital twin I/O system 40004. At 42314, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved vibration measurements and historical downtime data as inputs and calculates one or more outputs that represent probability of downtime values for machines in the manufacturing facility. Next, at 42318, the digital twin dynamic model system 40008 updates one or more probability of downtime values for machines in the manufacturing facility digital twins and all embedded digital twins based on the one or more outputs of the dynamic models.

[2183] FIG. 223 illustrates example embodiments of a method 42400 for updating one or more probability of shutdown values in the digital twin of an enterprise having a set of manufacturing facilities.

[2184] In the present example, the digital twin dynamic model system 40008 may receive requests from client application 40070 to update the probability of shutdown values for the set of manufacturing facilities within an enterprise digital twin. At 42402, digital twin dynamic model system 40008 receives a request from client application 40070 to update one or more probability of shutdown values of the enterprise digital twin and any embedded digital twins. Next, at 42404, digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore 40016. In this example, the digital twin dynamic model system 40008 may retrieve the digital twin of the enterprise and any embedded digital twins. At 42408, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore 400102. At 42410, the digital twin dynamic model system 40008 selects dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the and the one or more required inputs of the dynamic model(s) via

digital twin I/O system 40004. In the present example, the retrieved dynamic models may be configured to take one or more vibration measurements from vibration sensors 40036 and/or other suitable data as inputs and output probability of shutdown values for each manufacturing entity in the enterprise digital twin. At 42412, digital twin dynamic model system 40008 retrieves one or more vibration measurements from each of the selected vibration sensors 40036 from digital twin I/O system 40004. At 42414, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved vibration measurements and historical shut down data as inputs and calculates one or more outputs that represent probability of shutdown values for manufacturing facilities within the enterprise digital twin. Next, at 42418, the digital twin dynamic model system 40008 updates one or more probability of shutdown values of the enterprise digital twin and all embedded digital twins based on the one or more outputs of the dynamic model(s).

[2185] FIG. 224 illustrates example embodiments of a method 42500 for updating a set of cost of downtime values in machines in the digital twin of a manufacturing facility. In example embodiments, the manufacturing

[2186] In the present example, the digital twin dynamic model system 40008 may receive requests from a client application 40070 to populate real-time cost of downtime values associated with machines in a manufacturing facility digital twin. At 42502, digital twin dynamic model system 40008 receives a request from the client application 40070 to update one or more cost of downtime values of the manufacturing facility digital twin and any embedded digital twins (e.g., machines, machine parts, and the like) from the client application 40070. Next, at 42504, the digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins. In this example, the digital twin dynamic model system 40008 may retrieve the digital twins of the manufacturing facility, the machines, the machine parts, and any other embedded digital twins from digital twin datastore 40016. At 42508, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore 400102. At 42510, the digital twin dynamic model system 40008 selects dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system 40004. In the present example, the retrieved dynamic model(s) may be configured to take historical downtime data and operational data as inputs and output data representing cost of downtime per day for machines in the manufacturing facility. At 42512, digital twin dynamic model system 40008 retrieves historical downtime data and operational data from digital twin I/O system 40004. At 42514, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved data as input and calculates one or more outputs that represent cost of downtime per day for machines in the manufacturing facility. Next, at 42518, the digital twin dynamic model system 40008 updates one or more cost of downtime values of the manufacturing facility digital twins and machine digital twins based on the one or more outputs of the dynamic model(s).

[2187] FIG. 225 illustrates example embodiments of a method 42600 for updating a set of manufacturing KPI values in the digital twin of a manufacturing facility. In example embodiments, the manufacturing KPI is selected from the set of uptime, capacity utilization, on standard operating efficiency, overall operating efficiency, overall equipment effectiveness, machine downtime, unscheduled downtime, machine set up time, inventory turns, inventory accuracy, quality (e.g., percent defective), first pass yield, rework, scrap, failed audits, on-time delivery, customer returns, training hours, employee turnover, reportable health & safety incidents, revenue per employee, and profit per employee, schedule attainment, total cycle time, throughput, changeover time, yield, planned maintenance percentage, availability, and customer return rate.

[2188] In the present example, the digital twin dynamic model system 40008 may receive requests from a client application 40070 to populate real-time manufacturing KPI values in a manufacturing facility digital twin. At 42602, digital twin dynamic model system 40008 receives a request from the client application 40070 to update one or more KPI values of the manufacturing facility digital twin and any embedded digital twins (e.g., machines, machine parts, and the like) from the client application 40070. Next, at 42604, the digital twin dynamic model system 40008 determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins. In this example, the digital twin dynamic model system 40008 may retrieve the digital twins of the manufacturing facility, the machines, the machine parts, and any other embedded digital twins from digital twin datastore 40016. At 42608, digital twin dynamic model system 40008 determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore 400102. At 42610, the digital twin dynamic model system 40008 selects dynamic model input data sources (e.g., one or more sensors from sensor system 40030, data from Internet of Things connected devices 40024, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system 40030) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system 40004. In the present example, the retrieved dynamic model(s) may be configured to take one or more vibration measurements obtained from vibration sensors 40036 and other operational data as inputs and output one or more manufacturing KPIs for the facility. At 42612, digital twin dynamic model system 40008 retrieves one or more vibration measurements from each of the selected vibration sensors 40036 and operational data from digital twin I/O system 40004. At 42614, digital twin dynamic model system 40008 runs the dynamic model(s) using the retrieved vibration measurements and operational data as inputs and calculates one or more outputs that represent manufacturing KPIs for the manufacturing facility. Next, at 42618, the digital twin dynamic model system 40008 updates one or more KPI values of the manufacturing facility digital twins, machine digital twins, machine part digital twins, and all other embedded digital twins based on the one or more outputs of the dynamic model(s).

[2189] Further embodiments include a method for updating a set of biologically harmful agent concentration values in the digital twin of an industrial entity such as a wastewater treatment plant. Biologically harmful agents may be found in factories using metalworking fluids and may also be found in waste-handling facilities. Biologically harmful agents can be detected using biosensors.

In the present example, a client application, which interfaces with digital twin dynamic system, may be configured to provide a visualization of the concentration of a biologically harmful agent in the digital twin of the wastewater treatment plant. In example embodiments, biosensors may be acoustic biosensors, amperometric biosensors, electrochemical biosensors, optoelectric biosensors, calorimetric biosensors, potentiometric biosensors, immuno-biosensors, piezoelectric biosensors, and the like.

[2190] In this example, the digital twin dynamic system may receive requests from client application to update the biologically harmful agent concentration values in a wastewater treatment plant digital twin. At a next block, digital twin dynamic system receives a request from client application to update one or more biologically harmful agent concentration values of the wastewater treatment plant digital twin and any embedded digital twins from client application such that the concentration values represent real-time concentration levels of biologically harmful agents in the plant. At a next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In this example, the digital twin dynamic system may retrieve the digital twins of the wastewater treatment plant any other embedded digital twins. At a next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. At a next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. In the present example, the retrieved dynamic model(s) may be configured to take one or more concentration measurements obtained from biosensors, temperature measurements obtained from temperature sensors, and/or pressure measurements obtained from pressure sensors as inputs and output biologically harmful agent concentration measurements at different locations in the plant. At a next block, the digital twin dynamic system retrieves measurements from biosensors, temperature sensors, and/or pressure sensors disposed in the plant via digital twin I/O system. At a next block, digital twin dynamic system runs the dynamic model(s) using the retrieved measurements as inputs and calculates one or more outputs that represent biologically harmful agent concentration values at different locations in the wastewater treatment plant and/or throughout the plant. At a next block, the digital twin dynamic system updates one or more biologically harmful agent concentration values of the wastewater treatment plant digital twins, and all other embedded digital twins based on the output of the dynamic model(s).

[2191] Further example embodiments include a method for updating a set of fluid dynamics properties in the digital twin of an industrial entity such as a water supply piping system. In this example, a client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of the fluid flow rates in a water supply piping system in the digital twin of the water supply piping system. Fluid flow rates may depend on pressures, dimensions, and conduit material properties (shape, roughness, restrictions, and the like). Fluid

flow sensors may be configured to measure fluid flows. Fluid flow sensors may be flow meters, such as differential pressure flow meters (orifice plates, flow nozzles, Venturi tubes, variable area – rotameters), velocity flow meters, positive displacement flow meters, mass flow meters, and open channel flow meters (weirs, flumes, submerged orifices, current meters, acoustic flow meters, and the like).

[2192] In this example, the digital twin dynamic system may receive requests from client application to update the flow rate values in a water supply piping system digital twin. At the next block, digital twin dynamic system receives a request from client application to update one or more flow rate values in the piping system digital twin and any embedded digital twins such that the flow rate values represent real-time fluid flow rates in the piping system. At the next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In this example, the digital twin dynamic system may retrieve the digital twin of the water supply piping system, the digital twin of the facility containing the water supply piping system, and any other embedded digital twins. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. In the present example, the retrieved dynamic models may be configured to take one or more flow rate measurements obtained from the fluid flow sensors and model the flow rate values throughout the piping system. At the next block, digital twin dynamic system retrieves one or more measurements from each of the selected fluid flow sensors from digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved fluid flow rate measurements as inputs and calculates one or more outputs that represent flow rate values at different locations throughout the piping system and/or throughout the piping system. At the next block, the digital twin dynamic system updates one or more flow rate values of the water supply piping system digital twins, manufacturing facility digital twin, and all embedded digital twins based on the one or more outputs of the dynamic model(s).

[2193] Further example embodiments include a method for updating a set of radiation-related values in the digital twin of an industrial environment such as a nuclear production facility. Radiation modeling in a digital twin may be useful for nuclear energy production, nuclear research reactors, the fuel cycle, nuclear marine propulsion, and the like. Radiation sensors can use different types of detectors to measure site-specific levels of alpha, beta, gamma, or neutron radiation. In this example, client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of the gamma dose rate in the nuclear production facility in a digital twin of the nuclear production facility.

[2194] The digital twin dynamic system may receive requests from client application to update the gamma dose rate in the nuclear production facility digital twin. At the next block, digital twin dynamic system receives a request from client application to update one or more gamma dose rate values of the nuclear production facility digital twin and any embedded digital twins such that the gamma dose rates represent real-time gamma dose rates in the physical nuclear production facility system. At the next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In this example, the digital twin dynamic system may retrieve the digital twin of the nuclear production facility and any other embedded digital twins. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. In the present example, the retrieved dynamic models may be configured to take one or more gamma dose rate measurements obtained from the radiation sensors as inputs and output gamma dose rate values at other locations throughout the nuclear production facility. At the next block, digital twin dynamic system retrieves one or more measurements from each of the selected radiation sensors from digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved gamma dose rate measurements as inputs and calculates one or more outputs that represent gamma dose rate values at different locations in the facility and/or throughout the facility. At the next block, the digital twin dynamic system updates one or more gamma does rate values of the nuclear production facility digital twins and all embedded digital twins based on the one or more outputs of the dynamic model(s).

Quantum Systems for Industrial Internet of Things

[2195] Example embodiments include a method for updating a set of quantum mechanical values in the digital twin of an industrial environment. In this example, client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of quantum mechanical values in a digital twin of an industrial environment. For example, industrial entities that approach an atomic size will exhibit quantum mechanical behavior that may be modeled by dynamic models that adhere to quantum mechanical principles. Quantum mechanical properties may be measured by quantum sensor.

[2196] In this example, the digital twin dynamic system may receive requests from client application to update one or more quantum mechanical values of an industrial environment digital twin having embedded industrial entity digital twins representing industrial entities of an atomic size. At the next block, digital twin dynamic system receives a request from client application to update one or more quantum mechanical values of the industrial environment digital twin and the embedded digital twins such that the values represent real-time properties in the physical industrial environment. At the next block, digital twin dynamic system determines the one or more digital

twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In this example, the digital twin dynamic system may retrieve the digital twin of the industrial environment and the embedded atomic digital twins. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. In the present example, the retrieved dynamic model(s) may be configured take one or more quantum mechanical measurements obtained from quantum sensors disposed in the industrial environment as inputs and apply the one or more dynamic models, which adhere to quantum mechanics, to obtain one or more quantum mechanical values for different locations in the industrial environment and/or throughout the environment. At the next block, digital twin dynamic system retrieves one or more measurements from each of the selected quantum sensors via digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved quantum mechanical measurements as inputs and calculates one or more quantum mechanical values at different locations in the industrial environment and/or throughout the industrial environment. At the next block, the digital twin dynamic system updates one or more values of the industrial environment digital twin, atomic industrial entity digital twins, and all other embedded digital twins based on the one or more outputs of the quantum mechanical dynamic model(s).

[2197] Example embodiments include a method for updating a set of locations for an industrial entity such as a container in the digital twin of an industrial environment such as a manufacturing facility. In this example, client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of the location of containers through a manufacturing facility in the digital twin of the manufacturing facility.

[2198] In the present example, the digital twin dynamic system may receive requests from client application to update the locations values of the containers in a manufacturing facility digital twin. At the next block, digital twin dynamic system receives a request from client application to update one or more container location values in the manufacturing facility digital twin, embedded container digital twins, and any other embedded digital twins from client application such that the location values represent real-time locations of containers in the physical manufacturing facility. At the next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In this example, the digital twin dynamic system may retrieve the digital twin of the manufacturing facility, digital twins of the containers, digital twins of robots, and any other embedded digital twins. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more

sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic models using digital twin I/O system. In the present example, the retrieved dynamic models may adhere to classical dynamics. The one or more dynamic models may be configured take one or more velocity measurements obtained from Internet of Things connected devices used to move the containers, such as robots used to move the containers, as inputs and apply dynamic models to obtain one or more output values for container locations throughout the manufacturing facility. At the next block, digital twin dynamic system retrieves one or more velocity measurements from each of the selected robots via digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved velocity measurements as inputs and calculates one or more outputs that represent locations of containers throughout the environment. At the next block, the digital twin dynamic system updates one or more location values for containers of the manufacturing facility digital twin, container digital twins, robot digital twins, and all embedded digital twins based on the one or more outputs of the dynamic model(s).

[2199] Example embodiments include a method for updating a set of metal concentrations in an industrial environment such as a waste stream. In this example, client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of metal concentrations in a waste stream in a digital twin of the waste stream. For example, copper, chromium, nickel, and zinc are frequently found in high concentrations in industrial wastewater and each may be removed by precipitation.

[2200] The digital twin dynamic system may receive requests from client application to update the concentration of copper in an industrial waste stream digital twin. At the next block digital twin dynamic system receives a request from client application to update one or more copper concentration values of the waste stream digital twin and any other embedded digital twins (such as a precipitate filter digital twin) such that the copper concentration values represent real-time copper concentrations in the waste stream. At the next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In the present example, the digital twin dynamic system may retrieve the digital twin of the waste stream and any other embedded digital twins. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. In the present example, the retrieved dynamic models may adhere to inorganic chemistry principles. The dynamic models may take one or more copper concentration measurements obtained from chemical sensors disposed in the waste stream as inputs and apply dynamic models to obtain one or more outcome values for copper concentrations at

different locations in the waste stream and/or throughout the waste stream. At the next block, digital twin dynamic system retrieves one or more measurements from each of the selected chemical sensors via digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved measurements as inputs and calculates one or more outputs that represent the copper concentration values at different locations in the industrial waste stream and/or throughout the industrial waste stream. At the next block, the digital twin dynamic system updates one or more copper concentration values of the industrial waste stream digital twin and all embedded digital twins based on the one or more outputs of the dynamic model(s).

[2201] Example embodiments include a method for updating a set of organic compound concentrations in the digital twin of an industrial entity such as a container. In this example, client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of concentrations of an organic compound as in the digital twin of a container having a liquid and gas component.

[2202] In this example, the digital twin dynamic system may receive requests from client application to update the concentrations of the organic compound in a digital twin of a container. At the next block, digital twin dynamic system receives a request from client application to update one or more organic compound concentration values of the container digital twin and any other embedded digital twins such that the organic compound concentration values represent real-time organic compound concentrations in the container. At the next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins. In this example, the digital twin dynamic system may retrieve the digital twin of the container, digital twins that embed the container, and any other embedded digital twins from digital twin datastore. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. In the present example, the dynamic models may adhere to organic chemistry principles. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. The dynamic model(s) may be configured take one or more organic compound concentration measurements obtained from chemical sensors, temperature measurements from temperature sensor(s), and/or pressure measurements from pressure sensor(s) as inputs and apply dynamic models to obtain one or more output values for organic compound concentrations for different locations in the container and/or throughout the container. At the next block, digital twin dynamic system retrieves one or more measurements from each of the selected chemical sensors, temperature sensor, and pressure sensors via digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved measurements as inputs and calculates one or more outputs that represent the organic compound concentration values throughout the container. At the next block, the digital twin dynamic system updates one or more organic compound concentration

values of the container digital twin, all digital twins that embed the container, and all embedded digital twins based on the one or more outputs of the dynamic model(s).

[2203] Example embodiments include a method for updating a set of biological-related values in the digital twin of an industrial entity such as a beer brewing system. In this example, client application, which interfaces with the digital twin dynamic system, may be configured to provide a visualization of concentrations of a biological compound in the digital twin of a beer brewing system.

[2204] In this example, the digital twin dynamic system may receive requests from client application to update the concentrations of the biological compound in a digital twin of a beer brewing system. At the next block, digital twin dynamic system receives a request from client application to update one or more biological compound concentration values of the brewing system digital twin and any other embedded digital twins from such that the biological compound concentration values represent real-time concentrations in the physical process. At the next block, digital twin dynamic system determines the one or more digital twins required to fulfill the request and retrieves the one or more required digital twins from digital twin datastore. In this example, the digital twin dynamic system may retrieve the digital twin of the brewing system, digital twins of machine components, and/or any other embedded digital twins. At the next block, digital twin dynamic system determines one or more dynamic models required to fulfill the request and retrieves the one or more required dynamic models from dynamic model datastore. At the next block, the digital twin dynamic system selects dynamic model input data sources (e.g., one or more sensors from sensor system, data from Internet of Things connected devices, and any other suitable data) based on available data sources (e.g., available sensors from a set of sensors in sensor system) and the and the one or more required inputs of the dynamic model(s) via digital twin I/O system. In the present example, the retrieved dynamic models may adhere to biological principles.

[2205] The dynamic models may take one or more biological compound concentration measurements obtained from biosensors in the brewing system as inputs and apply dynamic models to obtain one or more output values for biological compound concentrations at different locations throughout the system. At the next block, digital twin dynamic system retrieves one or more measurements from each of the selected biosensors via digital twin I/O system. At the next block, digital twin dynamic system runs the dynamic model(s) using the retrieved biological compound concentration measurements as inputs and calculates one or more outputs that represent the biological compound concentration values at different locations in the system and/or throughout the system. At the next block, the digital twin dynamic system updates one or more biological compound concentration values of the brewing system digital twin and all embedded digital twins based on the one or more outputs of the dynamic model(s). In example embodiments, the digital twin dynamic system may be leveraged to enable a visual representation of a biological model in a digital twin of an industrial environment. In example embodiments, the biological model may be a biological population growth model. In example embodiments, the biological model may be a pathogen spreading model. In example embodiments, the biological model is an aging model.

[2206] FIG. 218 illustrates example embodiments of a display interface at 41200 that renders the digital twin of a dryer centrifuge, for example, and other information related to the dryer centrifuge. The display interface 41200 includes a title area at 41202 displaying any number of faults or other information related to the device. The display interface at 41200 can include a main screen at 41210 that can depict the machinery connections being monitored by the digital twin and rendered on the display interface 41200. The main screen 41210 can depict a left bearing 41302 connected to a motor 41304 having a right bearing 41308. The right bearing 41308 can be connected to a pulley 41340. The pulley 41340 can be connected to a belt 41350, which can be connected to a drive pulley 41360. The pulley 41360 can be connected to a left bearing 41370, which is connected to the dryer centrifuge 41372. The dryer centrifuge 41372 can have a right bearing 41374 connected to a pulley 413 8 zero. The pulley 41380 is connected to a belt 41390. The belt 41390 is connected to a pulley 41400. The pulley 41400 is connected to a left bearing 41410 of a motor 41412. The motor 41412 has a right bearing 41414. In these embodiments, motion of the left bearing 41302 can be depicted at 41320. Motion of the right bearing 41308 can be depicted at 41330. Motion of the left bearing 41370 can be depicted at 41420. Motion of the right bearing 41374 can be depicted at 41422. It will be appreciated in light of the disclosure that the display interface 41200 can be configured and reconfigured to display and depict motion (or characterizations of motion such as enlarged to more easily visualize) of one or more bearings and other machine components selected from the equipment available in the digital twin. The display interface 41200 further includes a detailed listing of each bearing or other relevant machine components at 41220 and the lifetime activity associated (or portions thereof) with those bearings. In example embodiments, such information can be inclusive of costs associated with the repair relevant to motion displayed by the digital twin. In example embodiments, these estimates can include a time to failure, a current probability to failure, cost of downtime, cost of repair, and the like. In example embodiments, display interface 41200 can depict the motion of the bearings and other relevant machine components at 41210 and can be depicted in a simplified graph at 41240 that can be selected between various locations at 41230 and can depict harmonic peaks at 41242, other relevant peaks 41244, and the like.

[2207] FIG. 226 illustrates example embodiments of a display interface at 45000 that renders the digital twin of a dryer centrifuge, for example, and other information related to the dryer centrifuge. The display interface 45000 includes a title area at 45002 displaying any number of faults or other information related to the device. The display interface at 41200 can include a main screen at 45010 that can depict the machinery connections being monitored by the digital twin and rendered on the display interface 41200. In this view, the user can adjust connections to depict the certain areas of the shop floor, manufacturing area, etc. where these machines can be located. In this view, the user can configure what is depicted on the main screen 45010 of the display interface 41200. In this view, not only can the user configure (and reconfigure) what is depicted on the main screen 45010 of the display interface 41200, the user can also configure (and reconfigure) to what connections the digital twin is listening and recording vibration, movement and other conditions at these connections. Further, the user can configure (and reconfigures) how the information received to

the display interface can be displayed. By way of these examples, the sensed information at 45020 can be configured (and reconfigured) to be displayed like the simplified motion by frequency at 41240 in FIG. 218, like what is shown at 46050 in FIG. 227.

[2208] FIG. 227 illustrates example embodiments of a display interface at 46000 that renders the digital twin of a dryer centrifuge, for example, and other information related to the dryer centrifuge. The display interface 46000 includes a title area at 46002 displaying any number of faults or other information related to the device. The display interface at 46000 can include a main screen at 46010 that can depict the machinery connections being monitored by the digital twin and rendered on the display interface 46000 similar to those in FIG. 346. The display interface 46000 further includes a detailed listing of each bearing and other relevant machine components at 46010 and the lifetime activity associated (or portions thereof) with those bearings. In example embodiments, such information can be inclusive of costs associated with the repair relevant to motion displayed by the digital twin. In example embodiments, these estimates can include a time to failure, a current probability to failure, cost of downtime, cost of repair, and the like. In example embodiments, display interface 41200 can depict the motion of the bearings and other relevant machine components at 46050 and can be depicted in a simplified graph at 46020 that can be selected between various locations at 46030 and can depict harmonic peaks at 46032, other relevant peaks 46034, filtered and combined views at 46042, and the like. The user can configure (and reconfigures) how the information received to the display interface can be displayed.

[2209] FIG. 228 illustrates example embodiments of a display interface at 47000 that renders the digital twin whose view at 47002 provides for selection between a digital twin dryer centrifuge at 47040, a digital twin lathe at 47010, a digital twin spinner at 47102, and the like. The digital twin dryer centrifuge 47040 includes the centrifuge and twin motor configuration at 47044 similar to what is depicted in FIG. 346. The digital twin dryer centrifuge 47040 can include a cost of repair indicator 47060 based on detected faults depicted at 47062. The digital twin dryer centrifuge 47040 can also include a cost of downtime indication at 47050 and a current probability of failure indicator at 47052. The digital twin lathe at 47010 can depict a motor 47012 connected to a lathe 47014. The digital twin lathe at 47010 can also include a cost of repair indicator 47030 based on detected faults depicted at 47032, a cost of downtime indication at 47020 and a current probability of failure indicator at 47022. Similar to the digital twin lathe at 47010, the digital twin spinner at 47102 can include a motor and spinner combination at 47100. As needed, the user can configure (and reconfigure) each of the views to add or modify what is being depicted.

[2210] FIG. 229 illustrates example embodiments of a display interface at 48000 that may render a digital twin whose view at 48002 incorporates connected machines each having drive bearings. The exemplary bearings 1, 2, 3, 4, 5, 6, 7, 8, 9 and 10 depicted in FIG. 229 can be displayed by the twin as two bearings 48012 and 48014 between a solid connection at 48010 correlating to bearings 1 and 2. Further, two bearings 48022 and 48024 between a solid connection at 48020 can correlate to bearings 3 and 4, two bearings 48032 and 48034 between a solid connection at 48030 can correlate to bearings 5 and 6, and so on. The display interface 48000 can include visualization controls at 48050 to control the view, the angles of the view and excitation frequencies. By way of

these examples, it can be seen that the two bearings 48032 and 48034 between the solid connection at 48030 are moving outside of nominal motion. In a sense, the user can plainly see what bearing or other component is ringing, vibrating, or otherwise moving outside of its nominal acceptable motion, which can be indicative of a need for repair, need for maintenance, and the like. It will also be appreciated in light of the disclosure that an issue causing the two bearings 48032 and 48034 between the solid connection at 48030 to vibrate can contribute to the motion of other bearings being outside of their nominal motion, as can be seen in bearing 7.

[2211] FIG. 230 illustrates example embodiments of a display interface at 48500 that may render a digital twin whose view at 48502 incorporates connected machines each having drive bearings. The exemplary bearings 1, 2, 3, and 4 depicted in FIG. 230 can be displayed by the twin as two bearings between a solid connection at 48520 correlating to bearings 1 and 2 and at 48530 correlating to bearings 3 and 4. By way of these examples, it can be seen that the two bearings and between the solid connection at 48520 (and to a lesser degree at 48530) are moving outside of nominal motion. Here too, the user can plainly see what bearing or other component is ringing, vibrating, or otherwise moving outside of its nominal acceptable motion, which can be indicative of a need for repair, need for maintenance, and the like. Here too, the two bearings and between the solid connection at 48520 vibrate and can contribute to the motion of other bearings being outside of their nominal motion, as can be seen in bearing 3 and 4. FIG. 231 illustrates example embodiments of a display interface at 48800 that may render a digital twin whose view at 48802 incorporates connected machines each having drive bearings like in FIG. 230. By way of these examples, it can be seen that the two bearings and between the solid connection at 48820 and 48830 are now moving nominally relative to what is shown in FIG. 230.

[2212] FIG. 232 illustrates example embodiments of a display interface at 49000 that may render a digital twin whose view at 49002 incorporates connected machines each having drive bearings. The exemplary bearings 1, 2, 3, and 4 can be displayed by the twin as two bearings 49022 and 49024 between a solid connection at 49020 correlating to bearings 1 and 2 and at bearings 49042 and 49044 between a solid connection at 49040 correlating to bearings 3 and 4. By way of these examples, it can be seen that the two bearings and between the solid connection at 49020 (and to a lesser degree at 49040) are moving outside of nominal motion. Here too, the user can plainly see what bearing or other component is ringing, vibrating, or otherwise moving outside of its nominal acceptable motion, which can be indicative of a need for repair, need for maintenance, and the like. Here too, the two bearings and between the solid connection at 49020 vibrate and can contribute to the motion of other bearings being outside of their nominal motion, as can be seen in bearing 3 and 4. Information about the motor and mill can be at 49060. In this example, the motor can drive the shaft from one end with a belt drive and such motion and one-sided drive can be noted in the view. FIG. 233 illustrates example embodiments of a display interface at 50000 that may render a digital twin whose view at 50002 incorporates connected machines each having drive bearings like in FIG. 232. By way of these examples, it can be seen that the two bearings 50012 and 50014 between the solid connection at 50010 and the two bearings 50022 and 50024 between the solid connection at 50020 are now moving nominally relative to what is shown in FIG. 232.

[2213] Referring to FIG. 234, the artificial intelligence system 55050 may define a machine learning model 55052 for performing analytics, simulation, decision making, and prediction making related to data processing, data analysis, simulation creation, and simulation analysis of one or more of the manufacturing entities 55010. The machine learning model 55052 is an algorithm and/or statistical model that performs specific tasks without using explicit instructions, relying instead on patterns and inference. The machine learning model 55052 builds one or more mathematical models based on training data to make predictions and/or decisions without being explicitly programmed to perform the specific tasks. The machine learning model 55052 may receive inputs of sensor data as training data, including event data 55140 and state data 55140 related to one or more of the manufacturing entities 55010. The sensor data input to the machine learning model 55052 may be used to train the machine learning model 55052 to perform the analytics, simulation, decision making, and prediction making relating to the data processing, data analysis, simulation creation, and simulation analysis of the one or more of the manufacturing entities 55010. The machine learning model 55052 may also use input data from a user or users of the information technology system. The machine learning model 55052 may include an artificial neural network, a decision tree, a support vector machine, a Bayesian network, a genetic algorithm, any other suitable form of machine learning model, or a combination thereof. The machine learning model 55052 may be configured to learn through supervised learning, unsupervised learning, reinforcement learning, self learning, feature learning, sparse dictionary learning, anomaly detection, association rules, a combination thereof, or any other suitable algorithm for learning.

[2214] The artificial intelligence system 55050 may also define the digital twin system 55070 to create a digital replica of one or more of the manufacturing entities 55010. The digital twin system 55070, the artificial intelligence system 55050, and the adaptive edge intelligence system 55060 can be included in the adaptive intelligence system 55080. The adaptive intelligence system 55080 can connect to the manufacturing entities 55010 through connectivity facilities 55020, which also permits connectivity with a monitoring system 55100 and a data collector system 55110. The digital replica of the one or more of the manufacturing entities may use substantially real-time sensor data to provide for substantially real-time virtual representation of the manufacturing entity and provides for simulation of one or more possible future states of the one or more manufacturing entities. The digital replica exists simultaneously with the one or more manufacturing entities 55010 being replicated. The digital replica provides one or more simulations of both physical elements and properties of the one or more manufacturing entities being replicated and the dynamics thereof. In example embodiments, throughout the lifestyle of the one or more manufacturing entities being replicated. The digital replica may provide a hypothetical simulation of the one or more manufacturing entities, for example during a design phase before the one or more manufacturing entities are constructed or fabricated, or during or after construction or fabrication of the one or more manufacturing entities by allowing for hypothetical extrapolation of sensor data to simulate a state of the one or more manufacturing entities, such as during high stress, after a period of time has passed during which component wear may be an issue, during maximum throughput operation, after one or more hypothetical or planned improvements have been made to

the one or more manufacturing entities, or any other suitable hypothetical situation. In example embodiments, the machine learning model 55052 may automatically predict hypothetical situations for simulation with the digital replica, such as by predicting possible improvements to the one or more manufacturing entities, predicting when one or more components of the one or more manufacturing entities may fail, and/or suggesting possible improvements to the one or more manufacturing entities, such as changes to timing settings, arrangement, components, or any other suitable change to the manufacturing entities. The digital replica allows for simulation of the one or more manufacturing entities during both design and operation phases of the one or more manufacturing entities, as well as simulation of hypothetical operation conditions and configurations of the one or more manufacturing entities. The digital replica allows for invaluable analysis and simulation of the one or more manufacturing entities, by facilitating observation and measurement of nearly any type of metric, including temperature, wear, light, vibration, etc. not only in, on, and around each component of the one or more manufacturing entities, but in some embodiments within the one or more manufacturing entities. In example embodiments, the machine learning model 55052 may process the sensor data including the event data 55140 and the state data 55130 from a data storage system 55120 to define simulation data for use by the digital twin system 55070. The machine learning model 55052 may, for example, receive state data 55130 and event data 55140 related to a particular manufacturing entity of the plurality of manufacturing entities and perform a series of operations on the state data 55130 and the event data 55140 to format the state data 55140 and the event data 55140 into a format suitable for use by the digital twin system 55070 in creation of a digital replica of the manufacturing entity. For example, one or more manufacturing entities may include a robot configured to augment products on an adjacent assembly line. The machine learning model 55052 may collect data from one or more sensors positioned on, near, in, and/or around the robot. The machine learning model 55052 may perform operations on the sensor data to process the sensor data into simulation data and output the simulation data to the digital twin system 55070. The digital twin system 55070 may use the simulation data to create one or more digital replicas of the robot, the simulation including for example metrics including temperature, wear, speed, rotation, and vibration of the robot and components thereof. The simulation may be a substantially real-time simulation, allowing for a human user of the information technology to view the simulation of the robot, metrics related thereto, and metrics related to components thereof, in substantially real time. The simulation may be a predictive or hypothetical situation, allowing for a human user of the information technology to view a predictive or hypothetical simulation of the robot, metrics related thereto, and metrics related to components thereof.

[2215] In example embodiments, the machine learning model 55052 and the digital twin system 55070 may process sensor data and create a digital replica of a set of manufacturing entities of the plurality of manufacturing entities to facilitate design, real-time simulation, predictive simulation, and/or hypothetical simulation of a related group of manufacturing entities. The digital replica of the set of manufacturing entities may use substantially real-time sensor data to provide for substantially real-time virtual representation of the set of manufacturing entities and provide for

simulation of one or more possible future states of the set of manufacturing entities. The digital replica exists simultaneously with the set of manufacturing entities being replicated. The digital replica provides one or more simulations of both physical elements and properties of the set of manufacturing entities being replicated and the dynamics thereof, in embodiments throughout the lifestyle of the set of manufacturing entities being replicated. The one or more simulations may include a visual simulation, such as a wire-frame virtual representation of the one or more manufacturing entities that may be viewable on a monitor, using an augmented reality (AR) apparatus, or using a virtual reality (VR) apparatus. The visual simulation may be able to be manipulated by a human user of the information technology system, such as zooming or highlighting components of the simulation and/or providing an exploded view of the one or more manufacturing entities. The digital replica may provide a hypothetical simulation of the set of manufacturing entities, for example during a design phase before the one or more manufacturing entities are constructed or fabricated, or during or after construction or fabrication of the one or more manufacturing entities by allowing for hypothetical extrapolation of sensor data to simulate a state of the set of manufacturing entities, such as during high stress, after a period of time has passed during which component wear may be an issue, during maximum throughput operation, after one or more hypothetical or planned improvements have been made to the set of manufacturing entities, or any other suitable hypothetical situation. In example embodiments, the machine learning model 55052 may automatically predict hypothetical situations for simulation with the digital replica, such as by predicting possible improvements to the set of manufacturing entities, predicting when one or more components of the set of manufacturing entities may fail, and/or suggesting possible improvements to the set of manufacturing entities, such as changes to timing settings, arrangement, components, or any other suitable change to the manufacturing entities. The digital replica allows for simulation of the set of manufacturing entities during both design and operation phases of the set of manufacturing entities, as well as simulation of hypothetical operation conditions and configurations of the set of manufacturing entities. The digital replica allows for invaluable analysis and simulation of the one or more manufacturing entities, by facilitating observation and measurement of nearly any type of metric, including temperature, wear, light, vibration, etc. not only in, on, and around each component of the set of manufacturing entities, but in some embodiments within the set of manufacturing entities. In example embodiments, the machine learning model 55052 may process the sensor data including the event data 55140 and the state data 55140 to define simulation data for use by the digital twin system 55070. The machine learning model 55052 may, for example, receive state data 55130 and event data 55140 related to a particular manufacturing entity of the plurality of manufacturing entities and perform a series of operations on the state data 55130 and the event data 55140 to format the state data 55140 and the event data 55140 into a format suitable for use by the digital twin system 55070 in the creation of a digital replica of the set of manufacturing entities. For example, a set of manufacturing entities may include a die machine configured to place products on a conveyor belt, the conveyor belt on which the die machine is configured to place the products, and a plurality of robots configured to add parts to the products as they move along the assembly

line. The machine learning model 55052 may collect data from one or more sensors positioned on, near, in, and/or around each of the die machines, the conveyor belt, and the plurality of robots. The machine learning model 55052 may perform operations on the sensor data to process the sensor data into simulation data and output the simulation data to the digital twin system 55070. The digital twin system 55070 may use the simulation data to create one or more digital replicas of the die machine, the conveyor belt, and the plurality of robots, the simulation including for example metrics including temperature, wear, speed, rotation, and vibration of the die machine, the conveyor belt, and the plurality of robots and components thereof. The simulation may be a substantially real-time simulation, allowing for a human user of the information technology to view the simulation of the die machine, the conveyor belt, and the plurality of robots, metrics related thereto, and metrics related to components thereof, in substantially real time. The simulation may be a predictive or hypothetical situation, allowing for a human user of the information technology to view a predictive or hypothetical simulation of the die machine, the conveyor belt, and the plurality of robots, metrics related thereto, and metrics related to components thereof.

[2216] In example embodiments, the machine learning model 55052 may prioritize collection of sensor data for use in digital replica simulations of one or more of the manufacturing entities. The machine learning model 55052 may use sensor data and user inputs to train, thereby learning which types of sensor data are most effective for creation of digital replicate simulations of one or more of the manufacturing entities. For example, the machine learning model 55052 may find that a particular manufacturing entity has dynamic properties such as component wear and throughput affected by temperature, humidity, and load. The machine learning model 55052 may, through machine learning, prioritize collection of sensor data related to temperature, humidity, and load, and may prioritize processing sensor data of the prioritized type into simulation data for output to the digital twin system 55070. In example embodiments, the machine learning model 55052 may suggest to a user of the information technology system that more and/or different sensors of the prioritized type be implemented in the information technology near and around the manufacturing entity being simulation such that more and/or better data of the prioritized type may be used in simulation of the manufacturing entity via the digital replica thereof.

[2217] In example embodiments, the machine learning model 55052 may be configured to learn to determine which types of sensor data are to be processed into simulation data for transmission to the digital twin system 55070 based on one or both of a modeling goal and a quality or type of sensor data. A modeling goal may be an objective set by a user of the information technology system or may be predicted or learned by the machine learning model 55052. Examples of modeling goals include creating a digital replica capable of showing dynamics of throughput on an assembly line, which may include collection, simulation, and modeling of, e.g., thermal, electrical power, component wear, and other metrics of a conveyor belt, an assembly machine, one or more products, and other components of the manufacturing ecosystem. The machine learning model 55052 may be configured to learn to determine which types of sensor data are necessary to be processed into simulation data for transmission to the digital twin system 55070 to achieve such a model. In example embodiments, the machine learning model 55052 may analyze which types

of sensor data are being collected, the quality and quantity of the sensor data being collected, and what the sensor data being collected represents, and may make decisions, predictions, analyses, and/or determinations related to which types of sensor data are and/or are not relevant to achieving the modeling goal and may make decisions, predictions, analyses, and/or determinations to prioritize, improve, and/or achieve the quality and quantity of sensor data being processed into simulation data for use by the digital twin system 55070 in achieving the modeling goal.

[2218] In example embodiments, a user of the information technology system may input a modeling goal into the machine learning model 55052. The machine learning model 55052 may learn to analyze training data to output suggestions to the user of the information technology system regarding which types of sensor data are most relevant to achieving the modeling goal, such as one or more types of sensors positioned in, on, or near a manufacturing entity or a plurality of manufacturing entities that is relevant to the achievement of the modeling goal is and/or are not sufficient for achieving the modeling goal, and how a different configuration of the types of sensors, such as by adding, removing, or repositioning sensors, may better facilitate achievement of the modeling goal by the machine learning model 55052 and the digital twin system 55070. In example embodiments, the machine learning model 55052 may automatically increase or decrease collection rates, processing, storage, sampling rates, bandwidth allocation, bitrates, and other attributes of sensor data collection to achieve or better achieve the modeling goal. In example embodiments, the machine learning model 55052 may make suggestions or predictions to a user of the information technology system related to increasing or decreasing collection rates, processing, storage, sampling rates, bandwidth allocation, bitrates, and other attributes of sensor data collection to achieve or better achieve the modeling goal. In example embodiments, the machine learning model 55052 may use sensor data, simulation data, previous, current, and/or future digital replica simulations of one or more manufacturing entities of the plurality of manufacturing entities to automatically create and/or propose modeling goals. In example embodiments, modeling goals automatically created by the machine learning model 55052 may be automatically implemented by the machine learning model 55052. In example embodiments, modeling goals automatically created by the machine learning model 55052 may be proposed to a user of the information technology system, and implemented only after acceptance and/or partial acceptance by the user, such as after modifications are made to the proposed modeling goal by the user.

[2219] In example embodiments, the user may input the one or more modeling goals, for example, by inputting one or more modeling commands to the information technology system. The one or more modeling commands may include, for example, a command for the machine learning model 55052 and the digital twin system 55070 to create a digital replica simulation of one manufacturing entity or a set of manufacturing entities, may include a command for the digital replica simulation to be one or more of a real-time simulation, and a hypothetical simulation. The modeling command may also include, for example, parameters for what types of sensor data should be used, sampling rates for the sensor data, and other parameters for the sensor data used in the one or more digital replica simulations. In example embodiments, the machine learning model 55052 may be

configured to predict modeling commands, such as by using previous modeling commands as training data. The machine learning model 55052 may propose predicted modeling commands to a user of the information technology system, for example, to facilitate simulation of one or more of the manufacturing entities that may be useful for the management of the manufacturing entities and/or to allow the user to easily identify potential issues with or possible improvements to the manufacturing entities.

[2220] In example embodiments, the machine learning model 55052 may be configured to evaluate a set of hypothetical simulations of one or more of the manufacturing entities. The set of hypothetical simulations may be created by the machine learning model 55052 and the digital twin system 55070 as a result of one or more modeling commands, as a result of one or more modeling goals, one or more modeling commands, by prediction by the machine learning model 55052, or a combination thereof. The machine learning model 55052 may evaluate the set of hypothetical simulations based on one or more metrics defined by the user, one or more metrics defined by the machine learning model 55052, or a combination thereof. In example embodiments, the machine learning model 55052 may evaluate each of the hypothetical simulations of the set of hypothetical simulations independently of one another. In example embodiments, the machine learning model 55052 may evaluate one or more of the hypothetical simulations of the set of hypothetical simulations in relation to one another, for example by ranking the hypothetical simulations or creating tiers of the hypothetical simulations based on one or more metrics.

[2221] In example embodiments, the machine learning model 55052 may include one or more model interpretability systems to facilitate human understanding of outputs of the machine learning model 55052, as well as information and insight related to cognition and processes of the machine learning model 55052, i.e., the one or more model interpretability systems allow for human understanding of not only "what" the machine learning model 55052 is outputting, but also "why" the machine learning model 55052 is outputting the outputs thereof, and what process led to the machine learning model 55052 formulating the outputs. The one or more model interpretability systems may also be used by a human user to improve and guide training of the machine learning model 55052, to help debug the machine learning model 55052, to help recognize bias in the machine learning model 55052. The one or more model interpretability systems may include one or more of linear regression, logistic regression, a generalized linear model (GLM), a generalized additive model (GAM), a decision tree, a decision rule, RuleFit, Naive Bayes Classifier, a K-nearest neighbors algorithm, a partial dependence plot, individual conditional expectation (ICE), an accumulated local effects (ALE) plot, feature interaction, permutation feature importance, a global surrogate model, a local surrogate (LIME) model, scoped rules, i.e., anchors, Shapley values, Shapley additive explanations (SHAP), feature visualization, network dissection, or any other suitable machine learning interpretability implementation. In example embodiments, the one or more model interpretability systems may include a model dataset visualization system. The model dataset visualization system is configured to automatically provide to a human user of the information technology system visual analysis related to distribution of values of the sensor data, the simulation data, and data nodes of the machine learning model 55052.

[2222] In example embodiments, the machine learning model 55052 may include and/or implement an embedded model interpretability system, such as a Bayesian case model (BCM) or glass box. The Bayesian case model uses Bayesian case-based reasoning, prototype classification, and clustering to facilitate human understanding of data such as the sensor data, the simulation data, and data nodes of the machine learning model 55052. In example embodiments, the model interpretability system may include and/or implement a glass box interpretability method, such as a Gaussian process, to facilitate human understanding of data such as the sensor data, the simulation data, and data nodes of the machine learning model 55052.

[2223] In example embodiments, the machine learning model 55052 may include and/or implement testing with concept activation vectors (TCAV). The TCAV allows the machine learning model 55052 to learn human-interpretable concepts, such as "running," "not running," "powered," "not powered," "robot," "human," "truck," or "ship" from examples by a process including defining the concept, determining concept activation vectors, and calculating directional derivatives. By learning human-interpretable concepts, objects, states, etc., TCAV may allow the machine learning model 55052 to output useful information related to the manufacturing entities and data collected therefrom in a format that is readily understood by a human user of the information technology system.

[2224] In example embodiments, the machine learning model 55052 may be and/or include an artificial neural network, e.g., a connectionist system configured to "learn" to perform tasks by considering examples and without being explicitly programmed with task-specific rules. The machine learning model 55052 may be based on a collection of connected units and/or nodes that may act like artificial neurons that may in some ways emulate neurons in a biological brain. The units and/or nodes may each have one or more connections to other units and/or nodes. The units and/or nodes may be configured to transmit information, e.g., one or more signals, to other units and/or nodes, process signals received from other units and/or nodes, and forward processed signals to other units and/or nodes. One or more of the units and/or nodes and connections therebetween may have one or more numerical "weights" assigned. The assigned weights may be configured to facilitate learning, i.e., training, of the machine learning model 55052. The weights assigned weights may increase and/or decrease one or more signals between one or more units and/or nodes, and in some embodiments may have one or more thresholds associated with one or more of the weights. The one or more thresholds may be configured such that a signal is only sent between one or more units and/or nodes, if a signal and/or aggregate signal crosses the threshold. In example embodiments, the units and/or nodes may be assigned to a plurality of layers, each of the layers having one or both of inputs and outputs. A first layer may be configured to receive training data, transform at least a portion of the training data, and transmit signals related to the training data and transformation thereof to a second layer. A final layer may be configured to output an estimate, conclusion, product, or other consequence of processing of one or more inputs by the machine learning model 55052. Each of the layers may perform one or more types of transformations, and one or more signals may pass through one or more of the layers one or more times. In example embodiments, the machine learning model 55052 may employ deep learning

and being at least partially modeled and/or configured as a deep neural network, a deep belief network, a recurrent neural network, and/or a convolutional neural network, such as by being configured to include one or more hidden layers.

[2225] In example embodiments, the machine learning model 55052 may be and/or include a decision tree, e.g., a tree-based predictive model configured to identify one or more observations and determine one or more conclusions based on an input. The observations may be modeled as one or more "branches" of the decision tree, and the conclusions may be modeled as one or more "leaves" of the decision tree. In example embodiments, the decision tree may be a classification tree. the classification tree may include one or more leaves representing one or more class labels, and one or more branches representing one or more conjunctions of features configured to lead to the class labels. In example embodiments, the decision tree may be a regression tree. The regression tree may be configured such that one or more target variables may take continuous values.

[2226] In example embodiments, the machine learning model 55052 may be and/or include a support vector machine, e.g., a set of related supervised learning methods configured for use in one or both of classification and regression-based modeling of data. The support vector machine may be configured to predict whether a new example falls into one or more categories, the one or more categories being configured during training of the support vector machine.

[2227] In example embodiments, the machine learning model 55052 may be configured to perform regression analysis to determine and/or estimate a relationship between one or more inputs and one or more features of the one or more inputs. Regression analysis may include linear regression, wherein the machine learning model 55052 may calculate a single line to best fit input data according to one or more mathematical criteria.

[2228] In example embodiments, inputs to the machine learning model 55052 (such as a regression model, Bayesian network, supervised model, or other type of model) may be tested, such as by using a set of testing data that is independent from the data set used for the creation and/or training of the machine learning model, such as to test the impact of various inputs to the accuracy of the model 55052. For example, inputs to the regression model may be removed, including single inputs, pairs of inputs, triplets, and the like, to determine whether the absence of inputs creates a material degradation of the success of the model 55052. This may assist with recognition of inputs that are in fact correlated (e.g., are linear combinations of the same underlying data), that are overlapping, or the like. Comparison of model success may help select among alternative input data sets that provide similar information, such as to identify the inputs (among several similar ones) that generate the least "noise" in the model, that provide the most impact on model effectiveness for the lowest cost, or the like. Thus, input variation and testing of the impact of input variation on model effectiveness may be used to prune or enhance model performance for any of the machine learning systems described throughout this disclosure.

[2229] In example embodiments, the machine learning model 55052 may be and/or include a Bayesian network. The Bayesian network may be a probabilistic graphical model configured to represent a set of random variables and conditional independence of the set of random variables. The Bayesian network may be configured to represent the random variables and conditional

independence via a directed acyclic graph. The Bayesian network may include one or both of a dynamic Bayesian network and an influence diagram.

[2230] In example embodiments, the machine learning model 55052 may be defined via supervised learning, i.e., one or more algorithms configured to build a mathematical model of a set of training data containing one or more inputs and desired outputs. The training data may consist of a set of training examples, each of the training examples having one or more inputs and desired outputs, i.e., a supervisory signal. Each of the training examples may be represented in the machine learning model 55052 by an array and/or a vector, i.e., a feature vector. The training data may be represented in the machine learning model 55052 by a matrix. The machine learning model 55052 may learn one or more functions via iterative optimization of an objective function, thereby learning to predict an output associated with new inputs. Once optimized, the objective function may provide the machine learning model 55052 with the ability to accurately determine an output for inputs other than inputs included in the training data. In example embodiments, the machine learning model 55052 may be defined via one or more supervised learning algorithms such as active learning, statistical classification, regression analysis, and similarity learning. Active learning may include interactively querying, by the machine learning model 55052, a user and/or an information source to label new data points with desired outputs. Statistical classification may include identifying, by the machine learning model 55052, to which a set of subcategories, i.e., subpopulations, a new observation belongs based on a training set of data containing observations having known categories. Regression analysis may include estimating, by the machine learning model 55052 relationships between a dependent variable, i.e., an outcome variable, and one or more independent variables, i.e., predictors, covariates, and/or features. Similarity learning may include learning, by the machine learning model 55052, from examples using a similarity function, the similarity function being designed to measure how similar or related two objects are.

[2231] In example embodiments, the machine learning model 55052 may be defined via unsupervised learning, i.e., one or more algorithms configured to build a mathematical model of a set of data containing only inputs by finding structure in the data such as grouping or clustering of data points. In example embodiments, the machine learning model 55052 may learn from test data, i.e., training data, that has not been labeled, classified, or categorized. The unsupervised learning algorithm may include identifying, by the machine learning model 55052, commonalities in the training data and learning by reacting based on the presence or absence of the identified commonalities in new pieces of data. In example embodiments, the machine learning model 55052 may generate one or more probability density functions. In example embodiments, the machine learning model 55052 may learn by performing cluster analysis, such as by assigning a set of observations into subsets, i.e., clusters, according to one or more predesignated criteria, such as according to a similarity metric of which internal compactness, separation, estimated density, and/or graph connectivity are factors.

[2232] In example embodiments, the machine learning model 55052 may be defined via semi-supervised learning, i.e., one or more algorithms using training data wherein some training examples may be missing training labels. The semi-supervised learning may be weakly supervised

learning, wherein the training labels may be noisy, limited, and/or imprecise. The noisy, limited, and/or imprecise training labels may be cheaper and/or less labor intensive to produce, thus allowing the machine learning model 55052 to train on a larger set of training data for less cost and/or labor.

[2233] In example embodiments, the machine learning model 55052 may be defined via reinforcement learning, such as one or more algorithms using dynamic programming techniques such that the machine learning model 55052 may train by taking actions in an environment in order to maximize a cumulative reward. In example embodiments, the training data is represented as a Markov Decision Process.

[2234] In example embodiments, the machine learning model 55052 may be defined via self-learning, wherein the machine learning model 55052 is configured to train using training data with no external rewards and no external teaching, such as by employing a Crossbar Adaptive Array (CAA). The CAA may compute decisions about actions and/or emotions about consequence situations in a crossbar fashion, thereby driving teaching of the machine learning model 55052 by interactions between cognition and emotion.

[2235] In example embodiments, the machine learning model 55052 may be defined via feature learning, i.e., one or more algorithms designed to discover increasingly accurate and/or apt representations of one or more inputs provided during training, e.g., training data. Feature learning may include training via principal component analysis and/or cluster analysis. Feature learning algorithms may include attempting, by the machine learning model 55052, to preserve input training data while also transforming the input training data such that the transformed input training data is useful. In example embodiments, the machine learning model 55052 may be configured to transform the input training data prior to performing one or more classifications and/or predictions of the input training data. Thus, the machine learning model 55052 may be configured to reconstruct input training data from one or more unknown data-generating distributions without necessarily conforming to implausible configurations of the input training data according to the distributions. In example embodiments, the feature learning algorithm may be performed by the machine learning model 55052 in a supervised, unsupervised, or semi-supervised manner.

[2236] In example embodiments, the machine learning model 55052 may be defined via anomaly detection, i.e., by identifying rare and/or outlier instances of one or more items, events and/or observations. The rare and/or outlier instances may be identified by the instances differing significantly from patterns and/or properties of a majority of the training data. Unsupervised anomaly detection may include detecting of anomalies, by the machine learning model 55052, in an unlabeled training data set under an assumption that a majority of the training data is "normal." Supervised anomaly detection may include training on a data set wherein at least a portion of the training data has been labeled as "normal" and/or "abnormal."

[2237] In example embodiments, the machine learning model 55052 may be defined via robot learning. Robot learning may include generation, by the machine learning model 55052, of one or more curricula, the curricula being sequences of learning experiences, and cumulatively acquiring new skills via exploration guided by the machine learning model 55052 and social interaction with

humans by the machine learning model 55052. Acquisition of new skills may be facilitated by one or more guidance mechanisms such as active learning, maturation, motor synergies, and/or imitation.

[2238] In example embodiments, the machine learning model 55052 can be defined via association rule learning. Association rule learning may include discovering relationships, by the machine learning model 55052, between variables in databases, in order to identify strong rules using some measure of "interestingness." Association rule learning may include identifying, learning, and/or evolving rules to store, manipulate and/or apply knowledge. The machine learning model 55052 may be configured to learn by identifying and/or utilizing a set of relational rules, the relational rules collectively representing knowledge captured by the machine learning model 55052. Association rule learning may include one or more of learning classifier systems, inductive logic programming, and artificial immune systems. Learning classifier systems are algorithms that may combine a discovery component, such as one or more genetic algorithms, with a learning component, such as one or more algorithms for supervised learning, reinforcement learning, or unsupervised learning. Inductive logic programming may include rule-learning, by the machine learning model 55052, using logic programming to represent one or more of input examples, background knowledge, and hypothesis determined by the machine learning model 55052 during training. The machine learning model 55052 may be configured to derive a hypothesized logic program entailing all positive examples given an encoding of known background knowledge and a set of examples represented as a logical database of facts.

[2239] In example embodiments, the platform can deploy many systems and methods for the industrial internet of things (IIoT) including solutions that can be configured as IIoT in a Box and other system configurations for IIoT; IIoT interface devices and systems (e.g., AR, VR, xR, wearables, and the like); advanced chips, boards, and switches for IIoT applications and the like. In example embodiments, the platform can deploy many different systems and methods for data collection, sensor fusion, data management and artificial intelligence; systems and methods for intelligent data collection for IIoT; systems and methods for equipment-specific data collection and management systems; systems and methods for biology-based data management for IIoT; systems and methods for advanced visual/optical sensing for IIoT intelligence; systems and methods for sensor fusion and sensor package configuration for IIoT intelligence; systems and methods for smart data pipelines for IIoT storage and computation; systems and methods for advanced, coordinated data collection and operations systems (e.g., drones, robotics, and the like); and systems and methods for advanced vibration sensing, monitoring and diagnostics. In example embodiments, the platform can deploy many systems and methods for advanced operational awareness and control including systems and methods for advanced industrial process control (e.g., hydrolysis to produce hydrogen for industrial heating, cooking, processing, etc.); systems and methods for artificial intelligence and data processing for detection and prediction of IIoT patterns and states; systems including platforms and associated methodologies for agile management and governance of IIoT operations (e.g., twins; dashboards; policy engine and the like); systems and methods for domain-specific applications of IIoT intelligence platform (e.g., oil & gas; mining,

agricultural, municipal, and the like); systems and methods for converged IIoT platforms; and systems and methods for automated industrial service ecosystems. In example embodiments, the platform can deploy many networking and computation for IIoT entities including systems and methods for convergence of edge and networking; systems and methods for enhancement of radio frequency (RF) networking for IIoT; systems and methods for quantum algorithms in combination with artificial intelligence for IIoT intelligence; and systems and methods for smart networking protocols.

[2240] The present disclosure relates to an enterprise and industrial control tower and enterprise management platform, referred to for simplicity herein as simply the enterprise management platform, or "EMP." In example embodiments, the enterprise management platform is configured to generate, integrate with, support, and/or operate on one or more digital twins. In general, digital twins merge data from multiple data sources into a model and representation of the salient characteristics of a thing, which is often a real-world physical thing, such as an operating environment or portion thereof (a building, campus, plant, factory, warehouse, distribution center, depot, port or the like), a set of items of equipment or infrastructure or components thereof, a product, or the like. A digital twin can also represent a process, such as a workflow, such as with moving elements that represent steps of the process, such as the flow of items through a plant or warehouse. A digital twin can also provide a logical representation, such as various topologies, clusters, networks, hierarchies or the like of logically related elements, such as an organizational chart of roles and/or personnel, the logical steps of a process, or the like. Thus, the term digital twin may refer to a digital representation of a thing or set of things. An enterprise digital twin may refer to any digital twin related to an enterprise and the wide array of things that relate to the enterprise and its operations. This may include digital twins of other enterprises and cohorts related to the enterprise, such as competitors, vendors, suppliers, distributors, customers, and the like. An enterprise may refer to a company, organization, corporation, LLC, non-profit organization, or the like. Enterprise digital twins may be used for a wide variety of user-facing applications that benefit from digital representation of salient features of elements of the enterprise, including monitoring of assets and operations, convenient generation and representation of a wide variety of analytic results, generation and display of simulations, such as for scenario planning, generation and display of recommendations and other decision support, collaborative decision support, and control of assets and operations, among many others. Enterprise digital twins may include organizational digital twins, executive digital twins, cohort digital twins, process digital twins, logical digital twins, real-time digital twins, AI-driven digital twins, environment digital twins, infrastructure and equipment digital twins, workforce digital twins, asset digital twins, product digital twins, system digital twins, and/or the like, which are discussed in greater detail throughout the disclosure.

[2241] In example embodiments, digital twins may be visual digital twins and/or data-based digital twins or combinations of visual and data-based digital twins. A visual digital twin may refer to a digital twin that is capable of being depicted in a display such as a traditional 2D display (optionally with touch, voice, optical, auditory, or other control features), a 3D display, an augmented reality display, a virtual-reality display, and/or a mixed-reality display, any of which may include various

combinations of computer-generated display elements (such as animations and other computer-generated graphics, including ones generated or derived from CAD and/or 3D models), elements captured by cameras (such as video and still images), visual elements captured or derived from various sensor systems, such as LIDAR and other point cloud systems, structured light systems, waveforms or other representations of information from acoustic sensor systems, vibration sensing systems, electromagnetic sensing systems, and many others, and/or elements captured, received, or derived from data collection and generation systems of enterprise assets, such as onboard diagnostic and reporting systems, IT systems (e.g., logs), information from wearable devices, and many others. A data-based digital twin may refer to a data structure that contains a set of parameters that are parameterized to represent a state of a thing or group of things. As used herein, the term "depict" may refer to the visual display of a thing and/or a digital representation of a thing in a data structure (e.g., in a data-based digital twin). It is noted that visual digital twins may also be data-based digital twins, or combinations of visual and data-based digital twins.

[2242] In example embodiments, a digital twin may be updated with real-time data, such that the digital twin reflects the state of a thing or set of things in real-time. For example, a digital twin of an operating environment or factory may depict the physical structure of the environment (e.g., walls, floors, ceilings, rooms and the like), as well as objects appearing in the environment (e.g., machines, products, employees, robots, and the like). Furthermore, depending on the manner in which this digital twin is configured, the digital twin of the operating factory may include things such as piping, conduits, wiring, foundations, and the like. In example embodiments, the digital twin may represent the information technology infrastructure of the factory, including wireless and fixed networking devices and systems and their operating capabilities and characteristics. In some implementations, the digital twin of the manufacturing environment may be updated with data received from sensors (e.g., IoT sensors deployed in or around a factory or equipment or machinery within the factory, wearable devices worn by workers within the factory, and other suitable data sources). For example, as a worker wearing a wearable device moves through the factory, the wearable device may communicate the relative location of the worker within the environment to the EMP, which in turn may update the digital twin to reflect the location of a representation of the worker in the digital twin of the factory. In scenarios where the digital twin is of a process, the digital twin may depict the process. For example, in the context of a manufacturing process, a digital twin of the process may depict the status and/or outcomes of different stages in the manufacturing pipeline. In some implementations, the EMP 60000 may receive data from various sources (e.g., IoT sensors, data from smart equipment, computing devices, smart products, smart infrastructure, or the like) and may update the digital twin of the process to reflect the received data. For example, in an automobile manufacturing factory, the process may include welding the chassis, placing the engine into the chassis, stamping the panels, assembling the body of the vehicle using the panels, painting the exterior, and assembling the interior of the vehicle. In this example, a process digital twin of the automobile assembly process may depict the various stages in the pipeline, and may further include metadata such as sensor readings, warnings, stoppages, delays, or the like. The EMP may be configured to generate, update, and/or provide enterprise digital twins

for different types of enterprises, including manufacturing enterprises, retail and marketing enterprises (merchants, advertisers, retail chains, restaurant chains, malls, and the like), technology enterprises (e.g., software, database and information technology companies), logistics enterprises (e.g., shipping and delivery entities), service-based enterprises (e.g., airlines, law firms, hospitals, accounting firms, and the like) and many others. For example, enterprise digital twins of a fast food enterprise may include digital twins of food production facilities, food production processes, food shipping facilities (e.g., warehouses and/or trucks), retail locations (e.g., individual restaurant locations), and/or retail processes (e.g., food preparation processes and/or customer workflows). In this example, these digital twins may identify the sources of contaminations (e.g., based on abnormal temperature readings in a food production factory), delays (e.g., based on outcomes of the production and/or shipping processes), customer satisfaction (e.g., based on data related to food preparation and/or customer workflows), and the like.

[2243] In example embodiments, the EMP may be configured to perform simulations using and/or with respect to one or more digital twins. In example embodiments, digital twins may be configured to behave in accordance with a set of constraints, such as laws of nature, laws of physics, mechanical properties, material properties, economic principles, chemical properties, and the like. In this way, the EMP may vary one or more parameters of a digital twin and may execute a simulation within the digital twin that conforms with real-word conditions and behaviors. For example, in executing a simulation of a logistics process that simulates outcomes associated with different packaging materials, the EMP may simulate variation of the packaging materials of one or more products. During the simulation, the products may be exposed to different conditions (e.g., different temperatures, humidity, motions, and the like). The simulation may be executed to determine the fraction of products that are likely to be damaged using the different packaging materials, which may affect the profitability of shipments vis-à-vis the cost of the different packaging materials and cost of replacing damaged products. In this way, the simulation may be run to help select the most cost-effective packaging material, such that estimated product loss is taken into account. Furthermore, in example embodiments, digital twins may be leveraged to perform simulations to predict future states of the thing or group of things and/or modeling behaviors in order to extrapolate states of the thing or group of things; to represent results of such simulations (including states, event and flows); and to offer opportunities to control things that are represented in the digital twins based on the simulations. For example, the EMP may receive sensor readings from temperature sensors, humidity sensors, and fan speed sensors deployed throughout an environment. The EMP may apply one or more thermodynamics equations to the received sensor readings and the dimensions of the environment to model the thermodynamic behavior of the environment to determine, to represent in the digital twin the temperatures in areas that do not have temperature sensors and to offer opportunities to adjust one or more systems, such as HVAC systems, or components thereof, to induce a change in the environment.

[2244] In example embodiments, the EMP is configured to generate organizational digital twins. In example embodiments, an organizational digital twin incorporates the organization chart ("org chart") of an enterprise. In example embodiments, an org chart may define the different divisions

(also referred to as business units) within an enterprise, the roles within each division, the reporting structure of the enterprise, and the individuals filling these roles. In example embodiments, the organizational digital twin may further include additional data for the business units, roles, and/or individuals filling the roles. For example, the organizational digital twin may include budgets for each business unit, salary ranges for roles, titles for roles, salaries for individuals, open roles, start dates for individuals, and the like. In example embodiments, an organizational digital twin may further incorporate data access rules for different divisions and/or roles within the organization, including permissions, access rights, and restrictions.

[2245] In example embodiments, an organizational digital twin may represent the organization as a hierarchy or other topology, where entities and relationships are represented, such as reporting relationships, relationships of authority or decision-making, or the like. In example embodiments, the organizational structure may be represented and maintained in a graph structure, such as a directed acyclic graph, a tree, or the like. In example embodiments, an organizational structure, such as an organizational chart or graph, may be parsed by an artificial intelligence system to automatically infer a set of entities, relationships, and roles, which in turn may be used to determine, or recommend, a set of default parameters for configuration of a digital twin. In example embodiments, the default parameters may be automatically configured for each user based on a role of the user within the organization, as inferred by the artificial intelligence system. In example embodiments, parameters may be adjusted by one or more authorized users, such as to adjust or correct the roles, using a digital twin configuration interface of the organizational digital twin. The parameters for configuration of a role-specific digital twin may include permissions (such as for data access), communication settings, availability of features (such as role-specific views of data and analytics, simulation features, control features, and many other features described throughout this disclosure), and the like. In example embodiments, the artificial intelligence services system may incorporate any of the techniques described throughout this disclosure or the documents incorporated by reference, such as a machine learning, deep learning, convolutional neural networks, robotic process automation, or the like. In example embodiments, the artificial intelligence system may include a machine learning system that is trained to infer roles within an organizational chart or structure based on a training set of data, such as one where roles and relationships within an organizational chart are provided by a set of human experts and/or where roles and relationships are explicitly stated within the organizational chart. For example, the artificial intelligence system may learn that the top of the organizational chart is likely to comprise the role of CEO and/or President of an organization, and that other roles, such as the CFO or COO, are likely to be represented in nodes that link directly to the CEO role. In example embodiments, the artificial intelligence system may be trained to operate on various data sources to determine and/or augment understanding of an organizational structure, such as public data sets, such as securities filings, social media information, web sites (such as securities information sites), public relations and other news about the organization, or the like. In example embodiments, the machine learning system may parse social media sites, such as LinkedIn□, to determine roles of individuals and/or to help infer roles. In example embodiments, data sources such as social data, web data,

news articles, or the like may be used to determine competencies of individuals, which may be associated with roles (e.g., the AI system may infer that a person with a finance degree is likely to be in a financial role within the organization). In example embodiments, settings for a user may be automatically configured to provide features that are appropriate for the training, education, experience and/or competencies of the user, as explicitly entered into the system or as inferred from information associated with the identity of the individual. For example, an individual who has a degree in physics and an MBA may be provided default access to physical model simulations and to financial simulations, while an individual who did not have those educational credentials might be required to obtain authorization and/or training before those features are made available in the digital twin. Thus, the EMP may include artificial intelligence systems that have been trained and/or configured to provide automated understanding of organizational structures and relationships, automated configuration of digital twins for roles within an organization based on the understanding of structures and relationships, and automated configuration of digital twin parameters, settings, and features based on the role and/or the identity of the user filling the role (including the competencies, education, experience, training, or the like of the user).

[2246] In example embodiments, a digital twin may be provided to represent the organizational structure of a third-party organization in the cohort of the organization of the user of the EMP, such as a supplier, vendor, distributor, logistics partner, value added reseller, representative, agent, venture partner, competitor, advertiser, marketplace or the like. An organizational digital twin of a cohort organization may represent structure, relationships, roles, identities, and competencies of individuals within roles or the organization, such that a user of the EMP may quickly and readily view salient information about the relevant parts of the organization. The organizational digital twin of the cohort organization may be automatically maintained by an artificial intelligence system of the EMP, such as by spidering, webscraping, and parsing websites, news feeds, press releases, social media data, and other available data sources, in order to maintain an accurate representation of the organization. The artificial intelligence system may be trained on a training set of data labeled by human users and/or automatically labeled to maintain an updated organizational structure. The resulting cohort digital twin may be configured to provide various role-specific views within the EMP. For example, a salesperson may be presented a digital twin view of the part of the cohort organization that is most likely to include individuals who are likely to be involved in a decision to purchase the user's offerings, while an HR person's view may be configured to present a digital twin view of the part of the cohort organization that provides the most comparable benchmark information for human resources. Digital twin views of cohort organizations may be automatically populated and/or configured, by training artificial intelligence systems on a process-specific or role-specific basis, to support a wide range of processes and features within the EMP, such as identification of recruiting candidates, benchmarking as to organizational structures, benchmarking as to competencies and talent, identification and/or configuration of sales and business development targets, identification of competitive offerings and/or projects, identification of targets for mergers and acquisitions, identification of targets for competitive research, and many others.

[2247] Digital twins can be helpful for visualizing the current state of a system, running simulations on those systems, and modeling behaviors, amongst many other uses. Depending on the configuration of the digital twin, however, a particular view or feature may not be useful for some members of an organization, as the configuration of the digital twin dictates the data that is depicted/visualized by the digital twin. Thus, as noted above, in example embodiments, the EMP is configured to generate role-based digital twins. Role-based digital twins may refer to digital twins of one or more segments/aspects of an enterprise, where the one or more segments/aspects and/or the granularity of the data represented by the role-based digital twin are tailored to a particular role within the entity and/or to the identity of a user that is associated with the role (optionally accounting for the competencies, training, education, experience, authority and/or permissions of the user, or other characteristics). In example embodiments, the role-based digital twins include executive digital twins. Executive digital twins may refer to digital twins that are configured for a respective executive within an enterprise. Examples of executive digital twins may include CEO digital twins, CFO (Financial) digital twins, COO (Operations) digital twins, HR digital twins, CTO (Technology) digital twins, CMO (Marketing) digital twins, General Counsel (Legal) digital twins, CIO (Information) digital twins, and the like. In some of these embodiments, the EMP generates different types of executive digital twins for users having different roles within the organization. In some of these embodiments, the respective configuration of each type of executive digital twin may be predefined with default digital twin data types, default relationships among entities, default features, and default granularities, among other elements. The default data types, entities, features and granularities may be determined based on a model of an organization, which may, in turn, be based on an industry-specific or domain-specific model or template, such as one that is based on a typical organizational structure for an industry (e.g., an automotive manufacturer, a consumer packaged goods manufacturer, a nationwide retailer, a regional grocery chain, or many others). In example embodiments, an artificial intelligence system may be trained, such as on a labeled industry-specific or domain-specific data set, to automatically generate an industry-specific or domain-specific digital twin for an instance of an EMP for an organization, with default configuration of data types, entities, features and granularities for various roles within an organization of that industry or domain. The defaults can then be reconfigured in a user interface of an authorized user to reflect company-specific variations from the industry-specific or domain-specific defaults. In example embodiments, a user (e.g., during an on-boarding process) may define the types of data depicted in the different types of executive digital twins, the entities to be represented, the features to be provided and/or the granularities of the different types of executive digital twins. Features may include what data is permitted to be accessed, what views are represented, levels of granularity of views, what analytic models and results can be accessed, what simulations can be undertaken, what changes can be made (including changes relevant to permissions of other users), communication and collaboration features (including receipt of alerts and the capacity to communicate directly to digital twins of other roles and users), control features, and many others. For convenience of reference, references to views, data, features, control or granularity throughout this disclosure should be understood to encompass

any and all of the above, except where context specifically indicates otherwise. Granularity may refer to the level of detail at which a particular type of data or types of data is/are represented in a digital twin. For example, a CEO digital twin may include P&L data for a particular time period but may not depict the various revenue streams and costs that contribute to the P&L data during the time period. Continuing this example, the CFO digital twin may depict the various revenue streams and costs during the time period in addition to the high-level P&L data. The foregoing examples are not intended to limit the scope of the disclosure. Additional examples and configurations of different executive digital twins are described throughout the disclosure.

[2248] In example embodiments, executive digital twins may allow a user (e.g., a CEO, CFO, COO, VP, Board member, GC, or the like) to increase the granularity of a particular state depicted in the digital twin (also referred to "drilling down into" a state of the digital twin). For example, a CEO digital twin may depict low granularity snapshots or summaries of P&L data, sales figures, customer satisfaction, employee satisfaction, and the like. A user (e.g., the CEO of an enterprise) may opt to drill down into the P&L data via a client application depicting the CEO digital twin. In response, the EMP may provide higher resolution P&L data, such as real-time revenue streams, real-time cost streams, and the like. In another example, the CEO digital twin may include visual indicators of different states of the enterprise. For example, the CEO digital twin may depict different colored icons to differentiate a condition (e.g., current and/or forecasted condition) of a respective data item. For example, a red icon may indicate a warning state, a yellow icon may indicate a neutral state, and a green icon may indicate a satisfactory state. In this example, the user (e.g., a CEO) may drill down into a particular data item (e.g., may select a red sales icon to drill down into the sales data, to see more specific and/or additional data, in order to determine why there is the warning state). In response, the CEO digital twin may depict one or more different data streams relating to the selected data item.

[2249] In example embodiments, a user interacting with a digital twin may escalate or deescalate a state to another user associated within an enterprise. For example, a COO or other operations executive may view a COO digital twin that depicts various operations related data. In this example, the COO may escalate a particular data set depicted in the COO digital twin to the CEO. Once escalated, the particular data set may appear in the CEO digital twin (e.g., with a message from the escalating executive).

[2250] In example embodiments, the EMP supports rolled-up real-time reporting. In some of these embodiments, data from IoT systems, sensors, onboard diagnostic systems, wearable devices, enterprise software systems, and/or other data sources (such as data feeds, news feeds, social media data sources, crowdsourced data, data obtained by spidering websites, sales data, marketing data, advertising data, market data, weather data, pricing data, and many others) may undergo one or more data fusion operations and an AI-based agent may determine which individuals within the enterprise to report results of analytics performed on the unfused or fused data. In example embodiments, the EMP may access data of or about an organization (and third-party or external data) that is available from a range of connected information technology and connectivity systems of the organization, including data collection, monitoring and storage systems as described

elsewhere in this disclosure and in the documents incorporated herein by reference. In example embodiments, the data collection, monitoring, and storage systems may include a "data pipeline" of such connected information technology and connectivity systems that may include one or more of individual sensors that are disposed on or around and/or are integrated into items (such as enterprise assets), packages of such sensors, data collection, detection and reading systems (such as asset tag readers, sensor readers, and many others); onboard diagnostic systems, log systems, and other onboard reporting systems producing feeds of data from machines, components or systems; networking devices, including switches, access points, routers, repeaters, mesh networking nodes, gateways, and the like, as well as a host of different types of smart or network-connected edge and IoT devices, and including BLUETOOTH™, BLUETOOTH LOW ENERGY™, WIFI, NFC, IR and other wireless devices, as well as 5G, 4G, 3G, LTE and other cellular infrastructure systems, including cellular chips and boards, gateways, towers and backhaul systems; data storage and processing systems, including local storage, distributed storage, database systems, caching systems, local memory systems, and many others; computational systems, including edge computational systems, serverless computational systems; and clients, servers, on-premises IT systems, cloud-based systems, and many others. Data may be transmitted and/or stored at points along this pipeline in raw form (such as in packets of raw data, with metadata, in streams, as events or transactions, as syndicated data, and in other forms) and/or in various processed forms, such as compressed data (including where compression is undertaken by trained artificial intelligence systems), summarized data (including where summarization is undertaking by trained artificial intelligence system), augmented data (such as by metadata and/or one or more analytic results), fused (e.g., multiplexed with one or more other sources), or the like. Collection, processing, storage and or transmission may be automated by one or more intelligence services systems as disclosed elsewhere in this document and the documents incorporated by reference herein, such as to provide for improved reliability, quality-of-service, efficiency, or the like, such as by intelligent protocol selection for data paths among nodes, intelligent filtering of RF-domain wireless transmission, and the like. As an example, a set of vibration sensors deployed on industrial machines/equipment in a factory may report vibration signatures of various components of the industrial machines/equipment. Edge devices may be configured to fuse the sensor data from an environment (e.g., a factory, warehouse, distribution center, office building, or many others) with other data collected with respect to the environment, whereby the fused data is fed to the digital twin. The EMP may then update the digital twin with the fused data and an AI system may analyze the digital twin and/or the fused data to identify data items to report, the proper role(s) to report to (e.g., CEO, COO, CFO, or the like), and then may provide the report to the appropriate individual(s). Enterprise digital twins, including executive digital twins, are discussed in greater detail throughout the application.

[2251] In example embodiments, the EMP may be configured to provide a set of collaboration tools that allow for collaboration between users (e.g., members of an organization and/or with third parties). In example embodiments, the collaboration tools allow users to collaborate with respect to and/or within one or more enterprise digital twins. In example embodiments, users can interact

while viewing the same digital twin or multiple digital twins showing different aspects of the enterprise, showing different views or features of the digital twin(s) and/or displaying information at different granularities.

[2252] In example embodiments, the collaboration tools include a video conferencing service. In some of these embodiments, the video conferencing service includes a graphical user interface that allows a user to create subchats during a video conference. A subchat may refer to an embedded video conferencing session where the members of the subchat are selected from an ongoing video chat. In example embodiments, the video conferencing service allows users to participate in video conferences within a digital twin. For example, users may access an environment digital twin via a VR-head set, whereby the participants may view the environment digital twin and see avatars of other participants within the "in-twin" video conference. In example embodiments, configuration of subchats may be created based on roles within an enterprise, such as where a role has authority to pull other roles into a subchat, such as roles that report to the authority role.

[2253] In example embodiments, the collaboration tools include interactive white boards, productivity tools (e.g., word processors, spreadsheets applications, slide decks/presentation applications, and the like), or some other type of collaboration tool. In these embodiments, users may import data from a digital twin (e.g., an executive twin) into a medium, such as into a word processor document or a spreadsheet. For example, when preparing a quarterly report, a CFO may import data from a CFO digital twin directly into the document containing the quarterly report. Collaboration tools are described in greater detail throughout the disclosure. In example embodiments, a digital twin may import data from one or more other collaboration environments into the digital twin, such that collaboration entities can be viewed alongside other entities represented in the digital twin. For example, a Google™ document containing an analytic report on the performance of a logistics system may be presented in a view of the elements of the logistics system in a digital twin.

[2254] In example embodiments, the EMP trains and deploys expert agents on behalf of enterprise users. In example embodiments, an expert agent is an AI-based software agent, using, for example, robotic process automation, that performs tasks on behalf of and/or suggests actions to a respective user having a defined role that requires some expertise in a particular field. In example embodiments, the expert agent may be trained within the EMP or otherwise, such as based on interactions of the user with a client application, such as actions taken by a user with respect to an executive digital twin, interactions with sensor data or other data collected by the EMP, interactions with systems or components of a workflow, and the like. In example embodiments, an expert agent may be an executive agent trained for executive roles. For example, an executive agent may be trained for performing or recommending actions to a user in an executive role, such as CEO role, a COO role, a CFO role, a CTO role, a CIO role, a CTO role, a CMO (chief marketing officer) role, a GC (general counsel) role, an HR (human resources) executive role, a board member role, a CDO (chief data officer) role, a CPO (chief product officer) role, and the like. In example embodiments, the EMP includes capabilities to train expert agents for other roles within an enterprise, such as an investor role, an engineering manager role, a project manager role, an

operations manager role, and a business development role, a factory manager role, a factory operations role, a factory worker role, a power plant manager role, a power plant operations role, a power plant worker role, an equipment service role, an equipment maintenance operator role, a logistics manager role, and the like.

[2255] In example embodiments, the expert agents are trained based on training data that includes actions taken by users and features relating to the circumstances surrounding the action (e.g., the type of action taken, the scenario that prompted the action, and the like). In example embodiments, the EMP receives telemetry data from a client application associated with a particular user and learns the workflows performed by the particular user based on the telemetry data and the surrounding circumstances. For example, the user may be a COO that is presented as a COO digital twin. Among the responsibilities of the COO may be scheduling maintenance and replacement of equipment or other infrastructure in a factory. The states depicted in the COO digital twin may include depictions of the condition of different pieces of equipment or infrastructure within the factory. In this example, the COO may schedule maintenance via the digital twin when a piece of equipment is determined to be in a first condition (e.g., a deteriorating condition) and may issue a request to the CFO via the COO digital twin for authorization of finances to replace the piece of equipment when the equipment is determined to be in a second condition (e.g., a critical condition). The executive agent may be trained to identify the COO's tendencies based on the COO's previous interaction with the COO digital twin. Once trained, the executive agent may automatically request replacements from the CEO when a particular piece of equipment is determined to be in the second condition and may automatically schedule maintenance if the piece of equipment is in the first condition. Further discussion of executive agents is provided throughout the disclosure. While reference is made to an expert agent being trained for a particular user, it is understood that an expert agent may be trained using the actions of one or more different users and may be used in connection with users that were not involved in training the expert agent. Expert agents, including executive agents, are discussed in greater detail throughout the disclosure.

[2256] FIG. 235 is a schematic of an example environment of an enterprise management platform 60000. In example embodiments, the EMP 60000 includes a configuration system 60002, a digital twin system 60004, a collaboration suite 60008, an expert agent system 60010, and an intelligence service system 60012. In example embodiments, the EMP 60000 includes an API system 60018 that facilitates the transfer of data between one or more external systems and the EMP 60000. In example embodiments, the EMP 60000 includes an enterprise data store 60014 that stores data relating to enterprises, whereby the enterprise data is used by the digital twin system 60004, the collaboration suite 60008, and/or the expert agent system 60010. The enterprise data store 60014 may store any of a wide variety of data, such as any data involved in the data pipeline described above and throughout this disclosure and the documents incorporated herein by reference. In example embodiments, the enterprise data store 60014 may store data that is being used to update digital twins in real-time or substantially real time. In example embodiments, the enterprise data store 60014 may store databases, file systems, folders, files, documents, transient data (e.g., real-time data or substantially real-time data), sensor data, and the like.

[2257] In example embodiments, the configuration system 60002 provides an interface (e.g., a graphical user interface (GUI)) by which a user (e.g., an "on-boarding" user) may upload or otherwise provide data relating to an enterprise. As used herein, an enterprise may refer to a for-profit or non-profit organization, company, governmental agency, non-governing organization, or the like. While described as an on-boarding user, the configuration of the enterprise management platform 60000 for a particular enterprise may be performed by any number of users, including individuals associated with the enterprise, individuals associated with the EMP, and/or individuals associated with a third-party, such as a third host of a hosted EMP for an enterprise (which may be deployed on cloud resources, platform-as-a-service, software-as-a-service, multi-tenant data resources and/or similar resources) and/or a service provider.

[2258] In example embodiments, the on-boarding user may define the types of enterprise digital twins that may be generated by the digital twin system 60004 on behalf of the enterprise being on-boarded. In example embodiments, the on-boarding user may select different types of digital twins that will be supported for the enterprise by the EMP 60000 via a GUI presented by the configuration system 60002. For example, the user may select different types of role-based digital twins from a menu of digital twin types, where the different types of role-based digital twins include executive digital twins. As another example, the user may select a type of organizational digital twin that is suitable for the user's organization, such as from a library of industry-specific or domain-specific organizational templates. In example embodiments, each type of executive digital twin has a predefined set of states (such term as referenced herein encompassing states, entities, relationships, parameters, and other characteristics) that are depicted in the respective executive digital twin and predefined granularity levels and/or other features for each state of the set. In example embodiments, the set of states that are depicted in the executive digital twin, the granularity of each, and/or other features may be customized (e.g., by the on-boarding user). In these embodiments, a user may define the different states that are represented in each type of executive digital twin and/or the granularity for each of the states depicted in the digital twin. For example, if the CEO of an enterprise has a financial background, the CEO may wish to have more financial data depicted in the CEO digital twin, such that the financial data is displayed at a higher granularity, or the CEO may wish to have access to underlying information on financial models that are available to the digital twin, such as models used for determination of state information (e.g., financial predictions or forecasts) or models used for augmentation of states (such as highlighting important deviations from expectations). By contrast, if the CEO has less financial experience or training, the CEO digital twin may be configured with summary financial data and may include prompts (which may be generated by an intelligent agent trained on a set of enterprise and/or industry outcomes) to obtain CFO input when states deviate from normal operating conditions. In this example, the CEO digital twin may be configured to depict the desired financial data fields at a granularity level set defined by a user (e.g., the financial data may include various revenue streams, cost streams, and the like). In another example, the CEO may have a technical background. In this example, the CEO digital twin may be configured to depict one or more states related to the enterprise's product and R&D efforts, patent development, and product roadmaps at

higher granularity levels. In yet another example, a COO may be tasked with overseeing a product team, a marketing team, and an HR department of the enterprise. In this example, the COO may wish to view marketing-related states, product development-related states, and HR-related states at a lower granularity level. In this example, the COO digital twin may be configured to show visual indicators that indicate whether any of the states are at a critical condition, an exceptional condition, or a satisfactory condition. For instance, if employee turnover is very high and employee satisfaction is low, the COO digital twin may depict that the HR-state is at a critical level. In this configuration, the COO may select to drill down into the HR-state, where she may view the employee turnover rate, hiring rate, and employee satisfaction survey results.

[2259] In another example, a COO or CTO digital twin may be configured to represent and assist with discovery and management of interconnections, relationships and dependencies between enterprise operations and information technology. For example, a COO digital twin or a CFO digital twin may be configured to depict a set of operations entities and workflows (e.g., flow diagrams that represent a production process, an assembly process, a logistics process, or the like), where entities (including human workers, robots, processing equipment, and other assets) are depicted to operate on a set of inputs such as materials, components, products, containers and information) in order to produce and hand off a set of outputs (of similar varied types) to the next set of entities in the workflow for further processing. These may be represented, for example, in a flow diagram that depicts each entity and its relationship in the flow to other entities. In example embodiments, a role-based digital twin (such as a CIO digital twin) may also represent an information technology system, such as representing sensors, IoT devices, data collection and monitoring systems, data storage systems, edge and other computational systems, wired and wireless networking systems, and the like, including any of the types described throughout this disclosure. Each information technology component or system may be depicted in the role-based digital twin, along with related data, such as specifications, configuration parameters and settings, processing capabilities, along with its relationship to other components, such as representing data and networking connectivity to other components or systems. In example embodiments, a role-based digital twin may provide a converged view that depicts operations technology entities and information technology entities in relation to each other, such as indicating which information technology entities are located with wired or proximal wireless connectivity to which operational entities, indicating which informational technology entities are logically associated to which operational entities (such as where cloud resources, computational resources, artificial intelligence resources, database resources, application resources, or other resources are provisioned to support or interact with operational entities, such as in virtual machine, container or other logical relationships). In example embodiments, the converged view presented in the role-based digital twin may thus depict location-based and/or logical interconnections between operations and information technologies. In example embodiments, alerts, such as indicating failure modes, congestion, delays, interruptions in service, poor latency, diminished quality of service, bandwidth constraints, poor performance on key performance indicators, downtime, or other issues may be provided as augmentations or overlays of the converged information technology and operations

digital twin, so that the COO, CTO, CIO or other users may see interconnections between information technology entities and operational entities that may be contributing to problems. Other types of issues that may be provided as augmentations or overlays may include alerts as to existing conditions and/or forecasts or predictions of such conditions, such as by analytic systems or forecasting artificial intelligence systems, such as expert agents trained to make such forecasts. In examples, if high latency in a control system for a warehouse is slowing down the process of picking and packing goods due to a related edge computational node experiencing congestion on an input data path, the user of the role-based digital twin may be alerted to the fact that operations are being adversely impacted by the congestion, and a recommendation may be presented to augment, update, upgrade, or replace either the system providing connectivity to the edge node or the edge node itself. Thus, a converged digital twin of operations technology entities and information technology entities may provide for insight into how an executive may adjust operations and/or information technology to improve results and/or avoid anticipated problems before they become catastrophic failures.

[2260] In example embodiments, a user (e.g., an on-boarding user) may connect one or more data sources 60030 to the EMP 60000. Examples of data sources 60030 that may be connected to the EMP may include, but are not limited to, a sensor system 60032 (e.g., a set of IoT sensors), a sales database 60034 that is updated with sales figures in real time, a customer relationship management (CRM) system 60038, a marketing campaign platform 60040, news websites, one or more financial databases 60048 that tracks costs of the business, surveys 60050 (e.g., customer satisfaction and/or employee satisfaction surveys), an Org Charts 60052, a workflow management system 60054, customer databases 60062 that store customer data, external data feeds (such as news feeds, public relations feeds, weather feeds, trade data, pricing data, market data, and the like), data obtained by spidering, webscraping, or otherwise parsing website and social media sites, data obtained by crowdsourcing, and/or data from many and various third-party datastores 60060 that store third-party data. The data sources 60030 may include additional or alternative data sources without departing from the scope of the disclosure. Once the user has defined the configuration of each respective executive digital twin, where the configuration includes the selected states to be depicted (which may include entities, relationships, and characteristics), the features that are to be enabled, and/or the desired granularity of each state, the user may then define the data sources 60030 that are fed into the respective executive digital twin, including any of the data sources in the data pipeline described above. In example embodiments, data from one or more of the data sources may be fused and/or analyzed before being fed into a respective digital twin.

[2261] In example embodiments, the user may define the data sources used to generate digital twins and to update the enterprise digital twins. In example embodiments, the user may define any physical locations that will be represented as an environment digital twin (which may be a digital twin of a factory or other suitable environments). For example, the user may define manufacturing facilities (e.g., factories), shipping facilities, warehouses, office buildings, and the like. Each factory may be given a location (which may include a logical and/or virtual location and/or a geo-location) and an identifier, such as a name and type description. In example embodiments, the

enterprise configuration system 60002 may assign an identifier to each factory and may associate the location of the factory with the identifier. In example embodiments, the user may define the types of objects that are included in the environment and/or may be found within an environment. For example, the user may define the types of enterprise resources (e.g., factory, warehouse, or distribution center equipment and machines, assembly lines, conveyors, vehicles, robots, high-lows, and the like, IT systems, workers, and many others) that are in the environment, the types of products, materials and components that are made in, stored in, moved around, assembled, used as inputs within, produced in, sold from, and/or received in the environment, the types of sensors/sensor kits and/or data collection, storage and/or processing devices that are used in the environment, the workers and workflows involved, and the like. Examples of how environment and process digital twins are generated and updated may be found herein.

[2262] In example embodiments, the configuration system 60002 is configured to generate organizational digital twins that represent an organizational structure of an enterprise. In example embodiments, the organizational digital twin may depict individuals/roles occupying the management and expert levels of an enterprise. Alternatively, the organizational digital twin may include a workforce digital twin that represents the entire workforce of an enterprise, including all the employees and/or contractors of the enterprise, or a defined part thereof. For example, in an enterprise setting, workforces may include a logistics workforce, a warehouse workforce, a distribution workforce, a reverse logistics workforce, a delivery workforce, a factory operations workforce, a plant operations workforce, a resource extraction operations workforce, a network operations workforce (e.g., for operating internal networks of an industrial enterprise), a sales workforce, a marketing workforce, an advertising workforce, a retail workforce, an R&D workforce, a technology workforce, an engineering workforce, and/or the like. In example embodiments, an organizational digital twin may include management-level roles within a workforce. Examples of management-level roles of an enterprise include a CEO role, a COO role, a CFO role, a counsel role, a board member role, a CTO role, an information technology manager role, a chief information officer role, a chief data officer role, an investor role, an engineering manager role, a project manager role, an operations manager role, a business development role. Furthermore, the management-level roles of a workforce may include a factory manager role, a factory operations role, a factory worker role, a power plant manager role, a power plant operations role, a power plant worker role, an equipment service role, and an equipment maintenance operator role. It is appreciated that not all of the roles defined above apply to a particular workforce type. Furthermore, some roles may be associated with different types of workforces.

[2263] In example embodiments, an organizational digital twin may further incorporate data access rules for different divisions and/or roles within the organization. For example, the CEO may be granted access to most or all of the organization's data, the CFO may be granted access to financial-related data and restricted from viewing R&D data, the CTO may be granted access to R&D-related data and restricted from viewing financial data, members of the engineering team may be restricted in accessing financial related data, or the like. Similar rules may be applied to access to features, such as analytic models, artificial intelligence systems, intelligent agents, and

the like, including role-based or identity-based control of the ability to view results, to configure inputs, to configure or adjust models (e.g., weights, inputs, or processing functions), to undertake control actions, or the like. In example embodiments, the EMP may utilize the organizational digital twin when determining the level of access a particular individual may be granted and/or whether to deny certain types of access to the individual. In example embodiments, the access rights may limit the types of data that particular users can access, such as information about each individual listed in the organizational digital twin (e.g., salary, start date, availability, work status, and the like). For example, lower level employees may not be granted access to sensitive information, such as financial data, product strategies, marketing strategies, trade secrets, or the like. In example embodiments, certain users may be granted permission to change the access rights of other employees, which may be reflected in the organizational digital twin. For example, certain executives and managers may be granted permission to grant access rights to members of their respective teams when working on certain projects.

[2264] In example embodiments, the configuration system 60002 receives an organization chart ("org chart") definition of an enterprise and generates an organizational digital twin based on the org chart definition. In example embodiments, the org chart definition may define the business units/departments of the enterprise, the reporting structure of the enterprise, various roles of the enterprise/within each business unit, and the individuals in the respective roles. In example embodiments, the user can upload the enterprise's org chart to the EMP 60000 via the enterprise configuration system 60002. Additionally or alternatively, the user can define the structure of the org chart (e.g., roles, business units, reporting structure) and may populate the various roles with names and/or other identifiers of the individuals filling the respective roles defined in the org chart. In example embodiments, the configuration system 60002 may access an enterprise resource planning system 60080 and/or an HR system 60082 of the enterprise to obtain organizational data of the enterprise, such as the roles of the enterprise, the individuals that fill the roles, the salaries of the individuals that fill the roles, the reporting structure of the enterprise, and the like. In these embodiments, the digital twin system 60004 (discussed below) may continue to communicate with the ERP system 60080 and/or HR system 60082 to receive the data needed to maintain the organizational digital twin in a real-time or near-real-time manner.

[2265] In example embodiments, the enterprise configuration system 60002 (in cooperation with the digital twin system 60004, discussed below) may generate an organizational digital twin of the enterprise based on the org chart definition and the individuals that populate the roles within the org chart definition. In example embodiments, a user may define one or more restrictions, permissions, and/or access rights of the individuals indicated in the organizational digital twin via the enterprise configuration system 60002. In example embodiments, a restriction may define one or more types of data or features that a particular user or group of users is not allowed to access (either directly or in a digital twin). In example embodiments, an access right may define one or more types of data or features that a particular user or group of users may access and the type of access that a user or group of users can access. In example embodiments, a permission may define operations that a user or a group of users may perform with respect to the EMP 60000. In example

embodiments, one or more of the access rights, permissions, and restrictions may be defined geographically and/or temporally limited. For example, some types of data or features may only be viewed or otherwise accessed in certain areas (e.g., sensitive data may only be viewed in the corporate offices) or at certain times (e.g., during Board meetings). In example embodiments, the restrictions, permissions, and/or access rights may be set with respect to roles or the users themselves. As such, defining access rights, permissions, and/or restrictions for a user or a group of users may also include defining access rights, permissions, and/or restrictions to a role and/or business unit within the enterprise. In example embodiments, the organizational digital twin may be deployed to manage the rights, permissions, and/or restrictions for the users of an enterprise. Furthermore, in example embodiments, the organizational digital twin may define the types of role-based digital twins (and other enterprise digital twins) that various users may have access to. In example embodiments, the organizational digital twin may depict additional or alternative information.

[2266] In example embodiments, the digital twin system 60004 is configured to generate, update, and serve enterprise digital twins of an enterprise. In example embodiments, the digital twin system 60004 is configured to generate and serve role-based digital twins on behalf of an enterprise and may serve the role-based digital twins to a client device 60102 (e.g., a mobile device, a tablet, a personal computer, a laptop, AR/VR-enabled device, workflow-specific device or equipment, or the like). As discussed, during the configuration phase, a user may define the different types of data and the corresponding data sources, data sets, and features that are used to generate and maintain each respective type of the different types of enterprise digital twins. Initially, the digital twin system 60004 configures the data structures that support each type of enterprise digital twin, including any underlying data sources/databases (e.g., SQL databases, graph databases, relational databases, distributed databases, blockchains, distributed ledgers, data feeds, data streams, and the like) that store or produce data that is ingested by the respective enterprise digital twins. Once the data structures that support a digital twin are configured, the digital twin system 60004 receives data from one or more data sources 60004. In example embodiments, the digital twin system 60004 may structure and/or store the received data in one or more databases. When a specific digital twin is requested (e.g., by a user via a client application 60104 or by a software component of the EMP 60000), the digital twin system may determine the views that are represented in the requested digital twin and may generate the requested digital twin based on data from the configured databases and/or real-time data received via an API. The digital twin system 60004 may serve the requested digital twin to the requestor (e.g., the client application or a backend software component of the EMP 60000). After an enterprise digital twin is served, some enterprise digital twins may be subsequently updated with real-time data received via the API system 60018. In example embodiments, an API may provide information to the data pipeline as to the type of data required for the digital twin, such that the data pipeline may be configured (by a user, or by an automated/intelligence system) to handle the data effectively. For example, the data pipeline may be configured to deliver data over a data path that uses an appropriate protocol for efficient delivery, delivering the data over a cost-appropriate path (e.g., an inexpensive path for data that

does not require low latency or real-time updating), or the like. Thus,, in example embodiments, configuration of a digital twin may include providing inputs as to the requirements of the digital twin for low-latency, high quality-of-service, high accuracy, high granularity, high reliability, or the like, based on, for example, the priority of the mission served by the data type. In example embodiments, an intelligent expert agent (or "intelligent agent" or "expert agent") may be trained on a training set of configurations of inputs to one or more data pipelines that were previously configured by experts, such that the intelligent agent may learn to automatically configure APIs for digital twins to provide appropriate inputs to data pipelines for subsequent digital twins involving similar or analogous workflows for similar or analogous roles, identities, industries and/or domains. In example embodiments, such training of an intelligent agent may include learning as to specific user interactions, such as learning which users within a role use which types of data at what times and for what purposes, such that data resources are appropriately allocated to support actual user requirements. For example, an automated intelligent agent managing the configuration of a data pipeline for a COO digital twin may learn that an operations executive (e.g., a COO user) checks production data for each factory at the end of each eight-hour shift (e.g., after 5:00 pm), such that mid-shift data updates are delivered over lower-cost data resources, but end-of-shift data is delivered over low-latency data paths that have high reliability and quality-of-service. Continuing this example, the intelligent agent may determine the frequency at which the production data is updated with respect to the COO digital twin, such that the COO digital twin is updated less frequently in the mornings and mid-afternoons, but is updated more frequently at the end of business hours. In example embodiments, the intelligent agent may be configured with business logic that defines overall strategies (e.g., when to use low-latency networks v. higher-latency networks and/or how often to update a certain type of data within a particular digital twin) and customized based on the preferences and use by the end user of the digital twin, whereby the overall strategies may be learned from training data sets obtained from experts and/or may be hard-coded by a developer, and the customization piece may be learned from monitoring the use of the digital twin by the end intended user (e.g., when she typically checks the production data of each factory). Additional or alternative examples of such data prioritization strategies and/or other configuration strategies should be understood to be encompassed herein. For example, upon receipt of inputs as to performance requirements, artificial intelligence capabilities of the data pipeline that is integrated with, linked to, or supporting of the EMP 100 may automatically or under user control employ techniques to provide appropriate resources at the right time and place, including, but not limited to: adaptive coding of data path transmissions between networked data communication nodes; adaptive filtering, repeating and amplification of RF/wireless signals (including software-implemented bandpass filtering); dynamic allocation of use of cellular and other wireless spectrum, adaptive, ad-hoc, cognitive management of wireless mesh network nodes; adaptive data storage; cost-based routing of wireless and wired signals; priority-based routing; channel- and performance-aware protocol selection for communications; context-aware allocation of computational resources, serverless computational systems, adaptive edge computational systems, channel-aware

error correction, smart-contract-implemented network resource allocation; and/or other suitable techniques.

[2267] In example embodiments, the digital twin system 60004 may be further configured to perform simulations and modeling with respect to the enterprise digital twins. In example embodiments, the digital twin system 60004 is configured to run data simulations and/or environment simulations using a digital twin. For example, a user may, via a client device, instruct the digital twin system 60004 to perform a simulation with respect to one or more states and/or workflows depicted in a digital twin. The digital twin system 60004 may run the simulation on the digital twin and may depict the results of the simulation in the digital twin. In this example, the digital twin may need to simulate at least some of the data used to run the simulation of the environment, so that there is reliable data when performing the requested environment simulation. The digital twin system 60004 is discussed in greater detail throughout the disclosure.

[2268] In example embodiments, the collaboration suite 60008 provides a set of various collaboration tools that may be leveraged by various users of an enterprise. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like. In example embodiments, an "in-twin" collaboration tool allows multiple users to view and collaborate within a digital twin. For example, in example embodiments, the collaboration tools may include an in-twin collaboration tool that enables a digital twin experience and a collaboration experience within the same interface (e.g., within an AR/VR-enabled user interface, a standard GUI, or the like), such as where collaboration entities and events (such as version-controlled objects, comment streams, editing events and other changes) are represented within the digital twin interface and linked to digital twin entities. For example, multiple users may be granted access to view an environment digital twin of a factory, such as a warehouse or factory, via an in-twin collaboration tool. Once viewing the environment digital twin, the users may then change one or more features of the environment depicted in the environment digital twin and may instruct the digital twin system to perform a simulation. In this example, the results of the simulation may be presented to the users in the digital twin and may be automatically populated into a shared document (e.g., a spreadsheet or presentation document). Users may collaborate in additional manners with respect to a digital twin, as will be discussed throughout the disclosure. For example, in example embodiments, the collaboration suite 60008 may allow a user to call a video conference with another user, where the users see each other and see aspects of a specific digital twin that relates to the topics of discussion for the conference. In this example, users may, for example, see a representation of a workpiece under discussion and see each other, so that a user can see gestures or indications from another user about how the workpiece should be acted upon. In another example, a conferencing feature of the twin may show participants in a view of a set of environments of facilities by their locations, so that users can recognize which participants may have the closest proximity to relevant assets that are the subject of the collaboration. In example embodiments, the collaboration suite 60008 interfaces with third-party applications, whereby data may be imported to and/or from the third-party application. For example, in collaborating on a Board presentation, different executives may export data from their

respective executive digital twin into a shared presentation file (e.g., PowerPoint™ file or Google™ slide presentation). In another example, a first user (e.g., the CEO of an enterprise) may request certain information (e.g., financial projections for the enterprise) from a second user (e.g., the CTO of the enterprise) via a first executive digital twin configured for the first user (e.g., a CEO digital twin of the enterprise). In response, the second user may upload/export the requested data from a second executive digital twin that was configured for the second user (e.g., the CTO) to the EMP 60000 (e.g., to the collaboration suite 60008 and/or the digital twin system 60004), which may then update the executive digital twin configured for the first user. Additional examples and descriptions of the collaboration suite 60008 and underlying collaboration tools are discussed throughout the disclosure.

[2269] In example embodiments, the collaboration suite 60008 may be configured to interface with the digital twin system 60004 (e.g., independent of or under control of the digital twin system 60004) to provide role-specific views and other features within a collaboration environment and/or workflow of a collaboration tool, such that different participants in the same collaboration environment and/or workflow experience different views or features of the same digital twin entities and/or workflows. For example, a CFO may collaborate with a COO and a CTO about the possible replacement of an internal system or a piece of machinery or equipment, where the current system, machinery or equipment and/or the potential replacement system, machinery, or equipment is/are represented in the digital twin by visual and other elements. During collaboration, the collaboration suite 60008 may recognize the identities/roles of the CFO, COO and CTO and may automatically configure their respective collaboration views into the example digital twin based on those roles. For example, the CFO may be presented with a view that is augmented with financial data, such as the cost of the item and various possible replacements, terms and conditions of leasing agreements, depreciation information, information on the financial impacts on productivity, or the like. Meanwhile, the collaboration suite 60008 may present the COO with information depicting the relationship of the item to operational processes, such as linkages to other systems involved in a production line, timing information (such as scheduled downtimes for a factory) and the like. In this example, the CTO may be presented with performance specifications and capability information for an item and various possible replacements, including, for example, compatibility information that indicates the extent to which various possible replacements are compatible with other items represented in the digital twin (including physical/mechanical compatibility, data compatibility, software compatibility, and many other forms of technology compatibility), reviews and ratings, and other technical information. Each executive user may be presented with respective information that is in the respective user's "native language" (e.g., information that is tailored to each executive's respective expertise and needs) and with respective views and/or features that are comfortable for that user, while the group can collaborate (in live or asynchronous modes) to raise issues, engage in commentary and dialog, perform analysis (including simulations as described herein) to arrive at a decision (e.g., about selection and timing of a replacement, or an alternative like a repair) that is financially prudent, operationally effective, and technologically sound. Thus, a role-sensitive collaboration environment integrated with respect to a shared enterprise digital

twin enables collaboration around digital twin entities and workflows while allowing users to engage with role-sensitive views and features. In example embodiments, the collaboration suite 60008 and or other systems of the EMP 60000 (e.g., the digital twin system 60004) may access a semantic model of an enterprise taxonomy to automatically generate and/or provide information that is presented in a shared digital twin (such as role-specific augmentation of entities with text or symbols that is derived from data or metadata based on state information or other data). In example embodiments, the enterprise taxonomy may be learned by the EMP 60000 via an analysis of data provided by the enterprise or may be manually uploaded by a user (e.g., a user configuration associated with the enterprise). The information in the digital twin may be presented with a role-specific understanding of the taxonomy, such as where the same entity (e.g., a piece of equipment) is given a different name by different groups in the enterprise (e.g., referred to as an "asset" by the finance department and a "machine" by the operations team) and/or where attributes of the entity or related workflows use different terminology, codes, symbols, or the like that are role-specific or group-specific. In example embodiments, the collaboration suite 60008 may automatically enable translation of terminology between roles, such as translating commentary that uses the name of an entity or that describes attributes of the entity from one role-specific form to another role-specific form. Automatic translation may present alternative terms together (e.g., as the "asset/machine" or "code red/urgent"). In example embodiments, automated translation may be performed by translation models (e.g., enterprise-specific translation models) that are trained by machine learning or similar techniques, whereby the translation models may be leveraged to provide automated translation for role-sensitive entity, workflow and attribute presentation. In example embodiments, the translation models may be trained using a training data set of translations generated by human experts and/or by unsupervised learning techniques that operate on the data of the enterprise to identify associations between different terms used by different roles and/or groups to describe the same thing. In example embodiments, translation models may be seeded by an explicit translation model or may be accomplished by deep learning or similar techniques known to those of skill in the art.

[2270] In example embodiments, the executive agent system 60010 trains expert agents that perform/recommend actions on behalf of an expert. An expert agent may be a software module that implements and/or leverages artificial intelligence services to perform/recommend actions on behalf of or in lieu of an expert. In example embodiments, an expert agent may include one or more machine-learned models (e.g., neural networks, prediction models, classification models, Bayesian models, Gaussian models, decision trees, random forests, and the like, including any of the artificial intelligence systems, expert systems, or the like described throughout this disclosure and/or the documents incorporated herein by reference) that perform machine-learning tasks, including robotic process automation, in connection with a defined role. Additionally or alternatively, an expert agent may be configured with artificial intelligence rules that determine actions in connection with a defined role. The artificial intelligence rules may be programmed by a user or may be generated by the expert agent system 60010. An expert agent may be executed at a client device 60102 and/or may be executed by or by a system that is linked to or integrated with

the EMP 60000. In example embodiments, the expert agent may be accessed as a service (e.g., via an API), such as in a service-oriented architecture, which in embodiments may be integrated with the EMP as service that is part of a microservices architecture. In embodiments where an expert agent is at least partially executed at a client device, the EMP 60000 may train an executive agent and may serve the trained executive agent to a client application 60104. In example embodiments, an expert agent may be implemented as a container (e.g., a Docker container), virtual machine, virtualized application, or the like that may execute at the client device 60102 or at the EMP 60000. In example embodiments, the executive agent is further configured to collect and report data to the expert agent system 60010, which the executive agent system 60010 uses to train/reinforce/reconfigure the expert agent. Many examples of such training are described throughout this disclosure and many others are intended to be encompassed by the disclosure.

[2271] In example embodiments, the executive agent system 60010 (working in connection with the artificial intelligence services system 60012) may train expert agents (e.g., executive agents and other expert agents), such as by using robotic process automation techniques, machine learning techniques, or other artificial intelligence or expert systems as described throughout this disclosure and/or the documents incorporated by reference herein to perform one or more executive actions on behalf of respective users, such as executives or other users who are responsible for undertaking activities that are automated by the robotic process automation or other techniques. In some of these embodiments, a client application 60104 may execute on a client device 60102 (e.g., a user device, such as a tablet, an AR and/or VR headset, a mobile device, or a laptop, an embedded device, an enterprise server, or the like) associated with a user (e.g., an executive, an administrative assistant of the executive, a board member, a role-based expert, a manager, a worker, or any other suitable employee or affiliate). In example embodiments, the client application 60104 may record the interactions of a user with the client application 60104 and may report the interactions to the executive agent system 60010. In these embodiments, the client application 60104 may further record and report features relating to the interaction, such as any stimuli or inputs that were presented to the user, what the user was viewing at the time of the interaction, the type of interaction, the role of the user, whether the interaction was requested by someone else, the role of the individual that requested the interaction, contextual information, state information, workflow information, event information, and the like. The executive agent system 60010 may receive the interaction data and related features and may generate, train, configure, and/or update an executive agent based thereon. In example embodiments, the interactions may be interactions by the user with an enterprise digital twin (e.g., an environment digital twin, a role-based digital twin, a process digital twin, and the like). In example embodiments, the interactions may be interactions by the user with data, such as sensor data (e.g., vibration data, temperature data, pressure data, humidity data, radiation data, electromagnetic radiation data, motion data, and/or the like) and/or data streams collected from physical entities of the enterprise (e.g., machinery, a building, a shipping container, or the like), data from various enterprise and/or third-party data sources (as described throughout this disclosure and incorporated documents), entity data (such as characteristics, features, parameters, settings, configurations, attributes and the like), workflow data (such as

timing, decision steps, events, tasks activities, dependencies, resources, or the like), and many other types of data. For example, a user may be presented with sensor data from a particular piece of machinery or equipment and, in response, may determine that a corrective action needs to be taken with respect to the piece of machinery or equipment. In this example, the expert agent may be trained on the conditions that cause the user to take a corrective action as well as instances where the user did not take corrective actions. In this example, the expert agent may learn the circumstances in which corrective action is taken.

[2272] In example embodiments, the executive agent system 60010 may train expert agents based on user interactions with network entities and/or computation entities. For example, the executive agent system 60010 may train an expert agent to learn the manner by which an IT expert diagnoses and handles a security breach. In this example, the expert agent may be trained to learn the steps undertaken by the expert to diagnose a security breach, the individuals within the enterprise that the security breach is reported to, and any actions undertaken by the expert to resolve the security breach.

[2273] In example embodiments, the types of actions that an expert agent may be trained to perform/recommend include: selection of a tool, selection of a task, selection of a dimension, setting of a parameter, configuration of settings, flagging an item for review, providing an alert, providing a summary report of data, selection of an object, selection of a workflow, triggering of a workflow, ordering of a process, ordering of a workflow, cessation of a workflow, selection of a data set, selection of a design choice, creation of a set of design choices, identification of a failure mode, identification of a fault, identification of an operating mode, identification of a problem, selection of a human resource, selection of a workforce resource, providing an instruction to a human resource, and providing an instruction to a workforce resource, amongst other possible types of actions. In example embodiments, an expert agent may be trained to perform other types of tasks, such as: determining an architecture for a system, reporting on a status, reporting on an event, reporting on a context, reporting on a condition, determining a model, configuring a model, populating a model, designing a system, designing a process, designing an apparatus, engineering a system, engineering a device, engineering a process, engineering a product, maintaining a system, maintaining a device, maintaining a process, maintaining a network, maintaining a computational resource, maintaining equipment, maintaining hardware, repairing a system, repairing a device, repairing a process, repairing a network, repairing a computational resource, repairing equipment, repairing hardware, assembling a system, assembling a device, assembling a process, assembling a network, assembling a computational resource, assembling equipment, assembling hardware, setting a price, physically securing a system, physically securing a device, physically securing a process, physically securing a network, physically securing a computational resource, physically securing equipment, physically securing hardware, cyber-securing a system, cyber-securing a device, cyber-securing a process, cyber-securing a network, cyber-securing a computational resource, cyber-securing equipment, cyber-securing hardware, detecting a threat, detecting a fault, tuning a system, tuning a device, tuning a process, tuning a network, tuning a computational resource, tuning equipment, tuning hardware, optimizing a system, optimizing a device, optimizing

a process, optimizing a network, optimizing a computational resource, optimizing equipment, optimizing hardware, monitoring a system, monitoring a device, monitoring a process, monitoring a network, monitoring a computational resource, monitoring equipment, monitoring hardware, configuring a system, configuring a device, configuring a process, configuring a network, configuring a computational resource, configuring equipment, and configuring hardware. As discussed, an expert agent is configured to determine an action and may output the action to a client application 60104. Examples of an output of an expert agent may include a recommendation, a classification, a prediction, a control instruction, an input selection, a protocol selection, a communication, an alert, a target selection for a communication, a data storage selection, a computational selection, a configuration, an event detection, a forecast, and the like. Furthermore, in example embodiments, the executive agent system 60010 may train expert agents to provide training and/or guidance rather in addition to or in lieu of outputting an action. In these embodiments, the training and/or guidance may be specific for a particular individual or role or may be used for other individuals.

[2274] In example embodiments, the executive agent system 60010 is configured to provide benefits to experts that participate in the training of expert agents. In example embodiments, the benefit is a reward that is provided based on the outcomes stemming from the user of an expert agent that is trained at least in part based on actions by the expert user. In example embodiments, the benefit is a reward that is provided based on the productivity of the expert agent. For example, if an expert agent trained by an individual is leveraged in connection with a set of users in the enterprise (or outside the enterprise), an account with the individual may be credited with a benefit such as cash rewards, stock rewards, gift card rewards, or the like. As the expert agent is used more, the benefit to the individual may be increased. In example embodiments, the benefit is a reward that is provided based on a measure of expertise of the expert agent. For example, individuals having a more sought after/valuable skill may be awarded greater benefits than individuals having a less sought after/valuable skill. In example embodiments, the benefit is a share of the revenue or profit generated by, or cost savings resulting from, the work produced by the expert agent. In example embodiments, the benefit is tracked using a distributed ledger (e.g., a blockchain) that captures information associated with a set of actions and events involving the expert agent. In some of these embodiments, a smart contract may govern the administration of the reward to the expert user.

[2275] In example embodiments, a set of expert agents trained by the executive agent system 60010 may be deployed as a double of at least a portion of a workforce of an enterprise, where the expert agents perform tasks of different roles within the enterprise. In some of these embodiments, the expert agents may be trained upon a training set of data that includes a set of interactions by members of a defined workforce of the enterprise during performance of the defined set of roles of the defined workforce (e.g., interactions with physical entities, digital twins, sensor data, data streams, computational entities, and/or network entities, among many others). In example embodiments, the interactions may be parsed to identify a chain of operations performed by the workforce and/or a chain of reasoning, whereby the chain of operations and/or chain of reasoning

are used to train the expert agents. In example embodiments, the interactions may be parsed to identify types of processing performed by the workforce upon a set of information, whereby the type of processing is embodied in the configuration of the respective expert agents. Examples of workforces may include, factory operations, plant operations, resource extraction operations, and the like.

[2276] In example embodiments, the executive agent system 60010 and/or a client application 60104 can monitor outcomes related to the user's interactions and may reinforce the training of the expert agent based on the outcomes. For example, each time the user takes a corrective action, the executive agent system 60010 may determine the outcome (e.g., whether a particular condition or issue was resolved) and whether the outcome is a positive outcome or a negative outcome. The executive agent system 60010 may then retrain the expert agent based on the outcome. Examples of outcomes may include data relating to at least one of a financial outcome, an operational outcome, a fault outcome, a success outcome, a performance indicator outcome, an output outcome, a consumption outcome, an energy utilization outcome, a resource utilization outcome, a cost outcome, a profit outcome, a revenue outcome, a sales outcome, and a production outcome. In these embodiments, the executive agent system 60010 may monitor data obtained from the various data sources after an action is taken to determine an outcome (e.g., sales increased/decreased and by how much, energy utilization decreased/increased and by how much, costs decreased/increased and by how much, revenue increased/decreased and by how much, whether consumption decreased/increased and by how much, whether a fault condition was resolved, and the like). The executive agent system 60010 may include the outcome in the training data set associated with the action undertaken by the expert that resulted in the outcome.

[2277] In example embodiments, the executive agent system 60010 receives feedback from users regarding respective executive agents. For example, in example embodiments, a client application 60104 that leverages an expert agent may provide an interface by which a user can provide feedback regarding an action output by an expert agent. In example embodiments, the user provides the feedback that identifies and characterizes any errors by the expert agent. In some of these embodiments, a report may be generated (e.g., by the client application or the EMP 60000) that indicates the set of errors encountered by the expert. The report may be used to reconfigure/retrain the executive agent. In example embodiments, the reconfiguring/retraining of an executive agent may include removing an input that is the source of the error, reconfiguring a set of nodes of the artificial intelligence system, reconfiguring a set of weights of the artificial intelligence system, reconfiguring a set of outputs of the artificial intelligence system, reconfiguring a processing flow within the artificial intelligence system, and/or augmenting the set of inputs to the artificial intelligence system.

[2278] In example embodiments, the expert agent may be configured to, at least partially, operate as a double of the expert for a defined role within an enterprise. In these embodiments, the executive agent system 60010 trains an expert agent based on a training data set that includes a set of interactions by a specific expert worker during the performance of their respective role. For example, the set of interactions that may be used to train the executive agent may include

interactions of the expert with the physical entities of an enterprise, interactions of the expert with an enterprise digital twin, interactions of the expert with sensor data obtained from a sensor system of the enterprise, interactions of the expert with data streams generated by the physical entities of the enterprise, interactions of the expert with the computational entities of the enterprise, interactions of the expert with the network entities, and the like. In example embodiments, the executive agent system 60010 parses the training data set of interactions to identify a chain of reasoning of the expert upon a set of interactions. In some of these embodiments, the chain of reasoning may be parsed to identify a type of reasoning of the worker, which may be used as a basis for configuring/training the expert agent. For example, the chain of reasoning may be a deductive chain of reasoning, an inductive chain of reasoning, a predictive chain of reasoning, a classification chain of reasoning, an iterative chain of reasoning, a trial-and-error chain of reasoning, a Bayesian chain of reasoning, a scientific method chain of reasoning, and the like. In example embodiments, the expert agent system parses the training data set of interactions to identify a type of processing undertaking by the expert in analyzing the set of interactions. For example, types of processing may include audio processing in analyzing audible information, tactile or "touch" processing in analyzing physical sensor information, olfactory processing in analyzing chemical sensing information, textual information processing in analyzing text, motion processing in analyzing motion information, taste processing in analyzing chemical information, mathematical processing in mathematically operating on numerical data, executive manager processing in making executive decisions, creative processing when deriving alternative options, analytic processing when selecting from a set of options, and the like.

[2279] In example embodiments, the expert agents include executive agents that are trained to output actions on behalf of executive and/or an administrator of an executive. In these embodiments, an expert agent may be trained for executive roles, such that a user in an executive role can train the executive agent by performing their respective role. For example, an executive agent may be trained for performing actions on behalf of or recommending actions to a user in an executive role. In some of these embodiments, the client application 60104 may provide the functionality of the enterprise management platform 60000. For example, in example embodiments, users may view executive digital twins and/or may use the collaboration tools via the client application 60104. During the use of the client application 60104, an executive may either escalate issues identified in the respective executive digital twin to another member of the enterprise. Each time the user interacts with the client application 60104, the client application 60104 may monitor the user's actions and may report the actions back to the expert agent system 60010. Over time, the executive agent system 60010 may learn how the particular user responds to certain situations. For instance, if the user is the CFO and each time a critical state with revenue or costs is identified in the CFO digital, the CFO escalates the critical state to the CEO, the executive agent system 60010 may learn to automatically escalate critical revenue states and critical cost states to the CEO. Further implementations of the executive agent system 60010 are discussed further in the disclosure.

[2280] In example embodiments, the artificial intelligence services system 60012 performs machine learning, artificial intelligence, and analytics tasks on behalf of the EMP 60000. In example embodiments, the artificial intelligence services system 60012 includes a machine learning system that trains machine learned models that are used by the various systems of the EMP 60000 to perform some intelligence tasks, including robotic process automation, predictions, classifications, natural language processing, and the like. In example embodiments, the EMP 60000 includes an artificial intelligence system that performs various AI tasks, such as automated decision making, robotic process automation, and the like. In example embodiments, the EMP 60000 includes an analytics system that performs different analytics across enterprise data to identify insights to various states of an enterprise. For example, in example embodiments, the analytics system may analyze the financial data of an enterprise to determine whether the enterprise is financially stable, in a critical condition, or a desirable condition. In example embodiments, the analytics system may perform the analytics in real-time as data is ingested from the various data sources to update one or more states of an enterprise digital twin. In example embodiments, the intelligence system includes a robotic process automation system that learns behaviors of respective users and automates one or more tasks on behalf of the users based on the learned behaviors. In some of these embodiments, the robotic process automation system may configure expert agents on behalf of an enterprise. The robotic process automation system may configure machine-learned models and/or AI logic that operate to output actions given stimulus. In example embodiments, the robotic process automation system receives training data sets of interactions by experts and configures the machine-learned models and/or AI logic based on the training data sets. In example embodiments, the artificial intelligence services system 60012 includes a natural language processing system that receives text/speech and determines a context of the text and/or generates text in response to a request to generate text. The intelligence services are discussed in greater detail throughout the disclosure.

[2281] In example embodiments, the EMP 60000 includes an enterprise data store 60014 that stores data on behalf of customer enterprises. In example embodiments, each customer enterprise may have an associated data lake that receives data from various data sources 60030. In example embodiments, the EMP 60000 receives the data via one or more APIs 60018. For example, in example embodiments, the API may be configured to obtain real-time sensor data from one or more sensor systems 60032 of an enterprise. The sensor data may be collected in a data lake associated with the enterprise. The digital twin system 60004 and the artificial intelligence services system 60012 may structure the data in the data lake and may populate one or more respective enterprise digital twins based on the collected data. In example embodiments, the data sources 60030 may include a set of edge devices 60064 that collect, receive and process data from a sensor system 60032, from suitable IoT devices, from local networking devices (e.g., wireless and fixed network resources, including repeaters, switches, mesh network nodes, routers, access points, gateways, and others), from general purpose networking devices (e.g., computers, laptops, tablets, smartphones and the like), from smart products, from telemetry systems of machinery, equipment, systems and components (e.g., onboard diagnostic systems, reporting systems, streaming systems,

syndication systems, event logs and the like), data collected by data collectors (including drones, mobile robots, RFID and other readers, and human-portable collectors) and/or other suitable data sources. In some of these embodiments, the edge devices 60064 may be configured to process sensor data (or other suitable data) collected at a "network edge" of the enterprise. Edge processing of enterprise data may include sensor fusion, data compression, computation, filtering, aggregation, multiplexing, selective switching, batching, packetization, streaming, summarization, fusion, fragmentation, encoding, decoding, transcoding, copying, storage, decompression, syndication, augmentation (e.g., by metadata), content inspection, classification, extraction, transformation, normalization, loading, formatting, error correction, data structuring, and/or many other processing actions. In example embodiments, the edge device 60064 may be configured to operate on the collected data and to adjust an output data stream or feed based on the contents of the collected data and/or based on contextual information, such as network conditions, operational conditions, environmental conditions, workflow conditions, entity state information, data characteristics, or many others. For example, an edge device 60064 may stream granular sensor data that is identified to be anomalous without compression, while the edge device 60064 may compress, summarize, or otherwise pass on a less granular data that is considered to be within a tolerance range of normal conditions or that reflects characteristics (e.g., statistical or signal characteristics) that suggest a lower likelihood that the data is likely to be of high interest. In this way, the edge device 60064 may provide semi-sentient data streams. Semi-sentience at the edge device 60064 may be improved by machine learning and training on a set of outcomes or feedback from users using process automation, machine learning, deep learning, or other artificial intelligence techniques as described herein. In example embodiments, the EMP 60000 may store the data streams in the data lake and/or may update one or more enterprise digital twins with some or all of the received data.

[2282] In example embodiments, the client devices 60102 may execute one or more client applications 60104 that interface with the EMP 60000. In example embodiments, a client application 60104 may request and display one or more enterprise digital twins. In some of these embodiments, a client application 60104 may depict an executive digital twin corresponding to the role of the user. For example, if the user is designated as the Chief Marketing Officer, the EMP 60000 may provide a CMO digital twin of the enterprise of the user. In some of these embodiments, the user data stored at the EMP 60000 and/or the client device 60102 may indicate the role of the user and/or the types of enterprise digital twins (and features thereof) to which the user has access.

[2283] In example embodiments, the client application 60104 may display the requested executive digital twin and may provide one or more options to perform one or more respective actions/operations corresponding to the executive digital twin and the states depicted therein. In example embodiments, the actions/operations may include one or more of "drilling down" into a particular state, escalating or otherwise notifying another user of a state or set of states, exporting a state or set of states into a collaborative environment (e.g., into a word processor document, a spreadsheet, a presentation document, a slide show, a model (e.g., a CAD model, a 3D model, or the like), a report (e.g., an annual report, a quarterly report, or the like), a website, a Wiki, a

dashboard, a collaboration environment location (e.g., a Slack location), a workflow application, or the like), sending a request for action with respect to one or more states from another user, performing a simulation, adjusting interface elements (such as changing sizes, colors, locations, brightness, presence/absence of display, etc.), or the like. For example, a COO or other operations executive may view an operations or COO digital twin. The states that may be depicted in the COO digital twin may include notifications of potential issues with one or more pieces of machinery or equipment (e.g., among many others, as observed from analyzing a stream of data from one or more sensors on a piece of robotic equipment). In viewing the COO digital twin, the user may wish to escalate the issue, such as to the CEO, request input from another executive and/or to instruct an operations manager, such as a warehouse or plant manager, to handle the issue. In this example, the client application depicting the COO digital twin may allow the user to select an option to escalate the issue. In response to the user selecting the "escalate" option, the client application 60104 transmits the escalate request to the EMP 60000. The EMP 60000 may then determine the appropriate user or users to which the issue is escalated. In example embodiments, the EMP 60000 may determine the reporting structure of the enterprise from an organizational digital twin of the enterprise to which the users belong. In this example, if the operations executive elects to have the operations manager handle the issue, the user may select an option to share the state with another user. The user may then enter an identifier of the intended recipient (e.g., an email address, phone number, text address, username, role description, or other identifiers of the recipient (such as identifiers for the recipient in various workflow environments, collaboration environments and the like (including other digital twins), and the like) and may input a message indicating instructions to the intended recipient. In response, the EMP 60000 may communicate the identified state to the intended recipient.

[2284] In another example, the client application 60104 may depict a CFO digital twin to a user (e.g., the CFO of an enterprise). In this example, the CFO may be tasked with preparing a quarterly report at the request of the CEO. In this example, the CFO may view a set of different financial states, including a P&L data, historical sales data (e.g., quarterly sales data and/or annual sales data), real-time sales data, projected sales data, historical cost data (e.g., quarterly costs and/or annual costs), projected costs, and the like. In this example, the CFO may select the states to include in the annual report, including the P&L data, quarterly sales data, and quarterly cost data. In response to the user selection, the client application 60104 may transmit a request to export the selected states into the annual report. In this example, the EMP 60000 may receive the request, identify the document (e.g., the annual report), and may include the selected states into the identified document.

[2285] In example embodiments, the client application 60104 may include a monitoring agent that monitors the manner by which a user responds to specific requests (e.g., a request from the CEO to populate a report) or notifications (e.g., a notification that a piece of machinery requires maintenance). The monitoring agent may report the user's response to such prompts to the EMP 60000. In response, the EMP 60000 may train an executive agent (which may include one or more machine-learned models) to handle such notifications when they next arrive. In example

embodiments, the monitoring agent may be incorporated in an executive agent that is incorporated in the client application 60104.

[2286] FIG. 237 illustrates an example set of components of the digital twin system 60004. In example embodiments, the digital twin system 60004 is executed by a computing system (e.g., one or more servers) that may include a processing system 60300 that includes one or more processors, a storage system 60330 that includes one or more computer-readable mediums, and a network interface 60350 that includes one or more communication units that communicate with a network (e.g., the Internet, a private network, and the like). In example embodiments, a processing system 60300 may execute one or more of a digital twin configuration system 60302, digital twin I/O system 60304, a data structuring system 60308, a digital twin generation system 60310, a digital twin perspective builder 60312, a digital twin access controller 60314, a digital twin interaction manager 60318, an environment simulation system 60320, a digital twin notification system 60322, and a digital twin simulation system 60320. The processing system 60300 may execute additional or alternative components without departing from the scope of the disclosure. In example embodiments, the storage system 60330 may store enterprise data, such as an enterprise data lake 60332, a digital twin data store 60334, a behavior datastore 60338 and/or other datastore, such as a distributed datastore, such as a set of blockchains or distributed data storage resources. The storage system 60330 may store additional or alternative data stores without departing from the scope of the disclosure. In example embodiments, the digital twin system 60004 may interface with the other components of the EMP 60000, such as the enterprise configuration system 60002, the collaboration suite 60008, the expert agent system 60010, and/or the artificial intelligence services system 60012.

[2287] In example embodiments, the digital twin configuration system 60302 is configured to set up and manage the enterprise digital twins and associated metadata of an enterprise, to configure the data structures and data listening threads that power the enterprise digital twins, and to configure features of the enterprise digital twins, including access features, processing features, automation features, reporting features, and the like, each of which may be affected by the type of enterprise digital twin (e.g., based on the role(s) that it serves, the entities it depicts, the workflows that it supports or enables and the like). In example embodiments, the digital twin configuration system 60302 receives the types of digital twins that will be supported for the enterprise, as well as the different states/objects/entities that will be depicted in each type of digital twin. For each type of digital twin, the digital twin configuration system 60302 identifies the types of data that feed or otherwise support each state/entity that is depicted in the respective type of digital twin and may determine any internal or external software requests that are required to support the identified data types. In example embodiments, the digital twin configuration system 60302 determines internal and/or external software requests that support the identified data types by analyzing the relationships between the different types of data that correspond to a particular state/entity/granularity/feature. In example embodiments, the digital twin configuration system 60302 determines and manages the data structures needed to support each type of digital twin. For an environment digital twin, for example, the digital twin configuration system 60302 may

instantiate a database (e.g., a graph database that defines the ontology of the environment and the objects existing (or potentially existing) within the environment and the relationships therebetween), whereby the instantiated database contains and/or references the underlying data that powers the environmental digital twin (e.g., sensor data and analytics, 3D maps, physical asset twins, and the like). In example embodiments, the different types of enterprise digital twins may be configured in accordance with a set of preference settings, taxonomy settings, topology settings, and the like. In example embodiments, the configuration system 60302 may utilize pre-defined preferences (e.g., default preference templates for different types of enterprise digital twins, including ones that are domain-specific, role-specific, industry-specific, workflow-specific and the like), taxonomies (e.g., default taxonomies for different types of enterprise digital twins), and/or topologies (e.g., default topologies for different types of twins, such as graph-based topologies, tree-based topologies, serial topologies, flow-based topologies, loop-based topologies, network-based topologies, mesh topologies, and others)) and/or receive custom preference settings and taxonomies from a configuring user. Examples of role-specific templates that are used to configure a role-based digital twin may include a CEO template, a COO template, a CFO template, a counsel template, a board member template, a CTO template, a chief marketing officer template, an information technology manager template, a chief information officer template, a chief data officer template, an investor template, a customer template, a vendor template, a supplier template, an engineering manager template, a project manager template, an operations manager template, a sales manager template, a salesperson template, a service manager template, a maintenance operator template, and/or a business development template. Similarly, examples of taxonomies that are used to configure different types of role-based digital twins may include CEO taxonomy, a COO taxonomy, a CFO taxonomy, a counsel taxonomy, a board member taxonomy, a CTO taxonomy, a chief marketing officer taxonomy, an information technology manager taxonomy, a chief information officer taxonomy, a chief data officer taxonomy, an investor taxonomy, a customer taxonomy, a vendor taxonomy, a supplier taxonomy, an engineering manager taxonomy, a project manager taxonomy, an operations manager taxonomy, a sales manager taxonomy, a salesperson taxonomy, a service manager taxonomy, a maintenance operator taxonomy, and/or a business development taxonomy.

[2288] In example embodiments, the digital twin configuration system 60302 may configure the databases that support each respective enterprise digital twin of an enterprise (e.g., role-based digital twins, environment digital twins, organizational digital twins, process digital twins, and the like), which may be stored on the digital twin data store 60334. In example embodiments, for each database configuration, the digital twin configuration system 60302 may identify and connect any external resources needed to collect data for each respective data type. For example, certain executive digital twins (e.g., CEO digital twin, CFO digital twin, COO digital twin, and CMO digital twin) may each require data derived and/or obtained from a CRM 60038 of the enterprise. In this example, the digital twin configuration system 60302 may configure one or more data collection threads to access an API, SDK, port, search facility, database access facility, and/or other connection facilities of the CRM 60038 of the enterprise on behalf of the enterprise and may obtain

any necessary security credentials to access the API. In another example, in order to collect data from one or more edge devices 60064 of the enterprise, the configuration system 60302 may initiate a process of granting access to the edge devices 60064 of the enterprise to the APIs of the EMP 60000.

[2289] In example embodiments, the digital twin I/O system 60304 is configured to obtain data from a set of data sources. In example embodiments, the digital twin I/O system 60304 (or other suitable components) may provide a graphical user interface that allows a user affiliated with an enterprise to upload various types of data that may be leveraged to generate the enterprise digital twins of the enterprise. For example, the user may upload 3D scans, still and video images, LIDAR scans, structured light scans, blueprints, 3D floor plans, object types (e.g., products, sensors, machinery, furniture, and the like), object properties (e.g., materials, physical properties, descriptions, price, and the like), output type (e.g., sensor units), architectural drawings, CAD documents, equipment specifications, and many others. In example embodiments, the digital twin I/O system 60304 may subscribe to or otherwise automatically receive data streams (e.g., publicly available data streams, such as RSS feeds, news streams, event streams, log streams, sensor system streams, and the like) on behalf of an enterprise. Additionally or alternatively, the digital twin system I/O system 60304 may periodically query and/or receive data from a connected data source 60030, such as a sensor system 60032 having sensors that sensor data from facilities (e.g., manufacturing facilities, shipping facilities, warehouse facilities, logistics facilities, retail facilities, distribution facilities, agricultural facilities, resource extraction facilities, computing facilities, transportation facilities, infrastructure facilities, networking facilities, data center facilities, and many others) and/or other physical entities of the enterprise, a sales database 60034 that is updated with sales figures in real time, a CRM system 60038, a content marketing platform 60040, an executive resource planning system 60034, financial databases 60048, surveys 60050, org charts 60052, workflow management systems 60054, third-party data sources 60060, customer databases 60062 that store customer data, and/or third-party datastores 60060 that store third-party data, edge devices 60064 that report data relating to physical assets (e.g., smart machinery/manufacturing equipment, sensor kits, autonomous vehicles, of the enterprise, wearable devices, and the like), enterprise management systems 60080, HR systems 60082, content management systems 60084, and the like). In example embodiments, the digital twin I/O system 60304 may employ a set of web crawlers to obtain data. In example embodiments, the digital twin I/O system 60304 may include listening threads that listen for new data from a respective data source.

[2290] In example embodiments, the digital twin I/O system 60304 is configured to serve the obtained data to instances of enterprise digital twins (which is used to populate digital twins) that are executed by a client device 60102 or the EMP 60000. In example embodiments, the digital twin I/O system 60304 receives data stream feeds received and/or collected on behalf of an enterprise and stores at least a portion of the streams into an enterprise data lake 60332 associated with the enterprise.

[2291] In example embodiments, the data structuring system 60308 processes data into a format that can be consumed by an enterprise digital twin. In example embodiments, processing by the

data structuring system 60308 may include compression, computation, filtering, aggregation, multiplexing, selective switching, batching, packetization, streaming, summarization, fusion, fragmentation, encoding, decoding, transcoding, encryption, decryption, duplication, deduplication, normalization, cleansing, identification, copying, storage, decompression, syndication, augmentation (e.g., by metadata), content inspection, classification, extraction, transformation, loading, formatting, error correction, data structuring, and/or many other processing actions. In example embodiments, the data structuring system 60308 may leverage ETL (extract, transform, load) tools, data streaming, and other data integration tooling to structure the data. In example embodiments, the data structuring system 60308 structures the data according to a digital twin data model that may be defined by the configuration system 60002 and/or a user. A data model may refer to an abstract model that organizes elements of enterprise-related data and standardizes the manner by which those elements relate to one another and to the properties of digital twin entities. For instance, a digital twin data model of an environment that includes vehicles (e.g., a vehicle assembly factory or an environment where vehicles operate) may specify that the data element representing a vehicle be composed of a number of other elements which represent sub-elements or attributes of the vehicle (the color of the vehicle, the dimensions of the vehicle, the engine of the vehicle, the engine parts of the vehicle, the owner of the vehicle, the performance specifications of the vehicle, and the like). In this example, the digital twin model components may define how the physical attributes are tied to respective physical locations on the vehicle. In example embodiments, digital twin models may define a formalization of the objects and relationships found in a particular application domain. For example, a digital twin model may represent the customers, products, and orders found in a manufacturing enterprise and how they relate to each other within the various digital twins. Additionally or alternatively, a digital twin data model may define a set of concepts (e.g., entities, attributes, relations, tables, and/or the like) used in defining such formalizations of data or metadata within the environment. For example, a "digital twin data model" used in connection with a banking application may be defined using the entity-relationship "data model" and how it is then related to the various executive digital twin views.

[2292] In example embodiments, the digital twin generation system 60310 serves enterprise digital twins on behalf of an enterprise. In some instances, the digital twin generation system 60310 receives a request for a specific type of digital twin from a client application 60104 being executed by a client device 60102 (e.g., via an API). Additionally or alternatively, the digital twin generation system 60310 receives a request for a specific type of digital twin from a component of EMP 60000 (e.g., the digital twin simulation system 60320). The request may indicate the enterprise, the type of digital twin, and the user (whose access rights may be verified or determined by the digital access controller 60314). In example embodiments, the digital twin generation system 60310 may determine and provide the client device 60102 with the data structures, metadata, ontology and information on hooks to data feeds as well as the digital twin constructs. This information may be used by the client to generate the digital twin in the end user device (e.g., an immersive device, such as AR devices or VR devices, tablet, personal computer, mobile, or the like). In example

embodiments, the digital twin generation system 60310 may determine the appropriate perspective for the requested digital twin (e.g., via the digital twin perspective builder 60312, which may include device-sensitive perspectives, such as delivering in appropriate formats based on the type of end user device) and any data restrictions that the user may have (e.g., via the digital twin access controller 60314). In response to determining the perspective and data restrictions, the digital twin generation system 60310 may generate the requested digital twin. In example embodiments, generating the requested digital twin may include identifying the appropriate data structure given the perspective and obtaining the data that parameterizes the digital twin, as well as any additional metadata that is served with the enterprise digital twin.

[2293] In example embodiments, the digital twin generation system 60310 may deliver the enterprise digital twin to the requesting client application 60104. In example embodiments, the digital twin generation system 60310 (or another suitable component) may continue to update a served digital twin with real-time data (or data that is derived from real-time data) as the real-time data is received and potentially analyzed, extrapolated, derived, predicted, and/or simulated by the EMP 60000.

[2294] In example embodiments, the digital twin generation system 60310 may obtain data streams from traditional data sources, such as relational databases, Hadoop file stores, graph databases that underlie operational and reporting tooling in the environment, telemetry data sources, onboard diagnostic systems, blockchains, distributed ledgers, distributed data sources, feed, streams, and many other sources. In example embodiments, the digital twin generation system 60310 may obtain data streams that are associated with the structural aspects of the data, such as the layout and 3D object properties of entities within facilities, the hierarchical design of a system of accounts, and/or the logical relationships of entities and actions in a workflow. In example embodiments, the data streams may include metadata streams that are associated with the nature of the data and data streams containing primary data (e.g., sensor data, sales data, survey data, and the like). For example, the metadata associated with a physical factory or other entity may include the types and layers of data that are being managed, while the primary data may include the instances of objects that fall within each layer. Layers for which metadata may be tracked and/or created may include, for example, metadata with respect to attributes, parameters or representations of a whole factory, component systems and assets within the factory (equipment, network entities, workforce entities, assets, and the like), sub-components and sub-systems, and further sub-components and sub-systems down to arbitrarily lower levels of granularity (e.g., a ball bearing of a rotating axle assembly of a fan that is part of a motor assembly driving an assembly line in a location of a warehouse). Layers may include, in another example, logical or operational layers, such as a reporting structure, such as from a COO to a VP of operations to a distribution manager to a warehouse manager to a shift manager to a warehouse worker. Layers may include workflow or process flow layers, such as from an overall process to its sub-components and decision points, such as an overall assembly process having sub-layers of gathering of input materials and components, positioning of workers, a series of assembly steps, inspection of outputs, and delivery to a post-assembly location.

[2295] In example embodiments, the digital twin perspective builder 60312 leverages metadata, artificial intelligence, and/or other data processing techniques to produce a definition of information required for generation of the digital twin in the digital twin generation system 60310. In example embodiments, different relevant datasets are hooked to a digital twin (e.g., an executive digital twin, an environment digital twin, or the like) at the appropriate level of granularity, thereby allowing for the structural aspects of the data (e.g., system of accounts, sensor readings, sales data, or the like) to be a part of the data analytics process. One aspect of making a perspective function is that the user can change the structural view or the granularity of data while potentially forecasting future events or changes to the structure to guide control of the area of the business at question. In example embodiments, the term "granularity" or "grain of data" may refer to a single line of data, a single byte of data, a single file, a single instance, or the like. Examples of "grains of data" or highly granular data may include a detailed record on a sale, a single block in a blockchain in a distributed ledger, a single event in an event log, a single vibration reading from a vibration sensor, or similar singular or atomic data units, and the like. Grain or atomicity may impose a constraint in how the data can be combined or processed to form different outputs. For example, if some element of data is captured only at the level of once-per-day, then it can only be broken down to single days (or aggregation of days) and cannot be broken down to hours or minutes, unless derived from the day representation (e.g., using inference techniques and/or models). Similarly, if data is provided only at the aggregate business unit level, it can be broken down to the level of an individual employee only by, for example, averaging, modeling, or inductive functions. Generally, role-based and other enterprise digital twins may often benefit from finer levels of data, as aggregations and other processing steps may produce outputs that are dynamic in nature and/or that relate to dynamic processes and/or real-time decision-making. It is noted that different types of digital twins may have different "sized" grains of data. For example, the grains of data that feed a CEO digital twin may be at a higher granularity level than the grains of data that feed a COO digital twin. In example embodiments, however, a CEO may drill down into a state of the CEO digital twin and the granularity for the selected state may be increased.

[2296] In example embodiments, the perspective builder 60312 adds relevant perspective to the data underlying the digital twin, which is provided to the digital twin generation system 60310. For example, a CEO digital twin may link in fuzzy data with markets data and depict the potential impacts of market forces on a simulated digital twin environment. In another example, in a CFO level digital twin, the internal financial system of accounts may be allocated across the physical structure of the digital twin providing an ability to understand the relationship between revenue generation, cost allocation, and the structural aspects of the business (e.g., the layout of a factory floor, a warehouse, a distribution center, a logistics facility, an office building, a retail location, a container ship, or the like). Continuing this example, the CTO digital twin may include data overlays with current market information on new technologies and linkages to outside information that may be used for enhancement of the factory. These different perspectives generated by the perspective builder 60312 combine with the digital twin generation system 60310 to give relevant simulations of how scenario-based future states might be handled by the factory, the digital twin

simulation system 60320 provides for recommendations on how to enhance the digital twin represented factory structurally to meet the needs of the future states. In example embodiments, the digital twin perspective builder 60312 may build perspectives that depict intersections or overlays of operational states and entities with information technology states and entities, which may facilitate recognition of opportunities and/or problems involving the interplay and convergence of information technology and operations technology within the operations of a wide range of industries and domains.

[2297] In example embodiments, a digital twin model is based on a combination of data and its relationship to the digital twin environments and/or processes. In example embodiments, different digital twins may share the same data and different digital twin perspectives can be the results of a set of metadata built on top of a digital twin data model or data environment. In example embodiments, the digital twin data model provides the details of the information to be stored and it is used to build a layered system where the final computer software code is able to represent the information in the lower levels in a form that is appropriate for the digital twin perspective being used. FIG. 238 illustrates a business model of an enterprise as a combination of two parallel design and implementation hierarchies. The left-hand side of FIG. 238, at 60504, illustrates the manner by which processes that are to be represented in the digital twin are designed and the data flows are set up to direct these to the application programs. In this example, the right-hand side (the data model side), at 60502, illustrates the manner by which digital twin data models are designed and deployed to allow the application programs to be fed with either batch or real-time data.

[2298] In example embodiments, the digital twin access controller 60314 informs the digital twin generation system 60310 of specific constraints around the roles of users able to view the digital twin as well as providing for dynamically adjustable digital twins that can adapt to constrain or release views of the data or other features specific to each user role. For example, sensitive salary data might be obfuscated from most administrative employees when viewing an organizational digital twin, but the CEO may be granted access to view the salary information directly. In example embodiments, the digital twin access controller 60314 may receive a user identifier and one or more data types. In response, the digital twin access controller 60314 may determine whether the user indicated by the user identifier has access to the one more data types or other features. In some of these embodiments, the digital twin access controller may look up the user in the organizational digital twin of the enterprise of the user and may determine the user's permissions and restrictions based thereon. Alternatively, the user's permissions and restrictions may be indicated in a user database. In example embodiments, the organizational digital twin may, as noted above, be generated automatically, such as by parsing available data sources to automatically construct a representation of the organization, such as a hierarchical organizational chart, a graph of the organization with nodes representing organizational entities (e.g., workgroups, roles, assets and personnel), links or connections indicating relationships (e.g., reporting relationships, lines of authority, group affiliations, and the like), and data or metadata indicating other attributes of the entities and relationship, and the like.

[2299] In example embodiments, the digital twin interaction manager 60318 manages the relationship between the structural view of the data in an enterprise digital twin (e.g., as depicted/represented by the client application 60104) and the underlying data streams and data sources. In example embodiments, this interaction layer makes the digital twin into a window into the underlying data streams through the lens of the structure of the data. In example embodiments, the digital twin interaction manager 60318 determines the types of data that are being fed to an instance of an enterprise digital twin (e.g., an environment digital twin or an executive digital twin) while the instance is being executed by a client application 60104. Put another way, the digital twin interaction manager 60318 determines and serves data for an in-use digital twin. In example embodiments, the digital twin interaction manager 60318 feeds raw data received from a data source to the digital twin. For example, sensor readings of temperatures throughout an environment may be fed directly to the executing environment digital twin of the environment. In example embodiments, the digital twin interaction manager 60318 obtains data and/or instructions that are derived by another component of the EMP 60000. For example, the digital twin interaction manager 60318 may obtain analytical data from the artificial intelligence services system 60012 that is derived from incoming financial data, marketing data, operational data, and sensor data. In this example, the digital twin interaction manager 60318 may then feed the analytical data to an executive digital twin (e.g., CEO digital twin), whereby the analytical data may be conveyed to the user. In another example, the digital twin interaction manager 60318 may receive simulated cost data from the digital twin simulation system 60320 to convey revenue/costs with respect to different asset maintenance schedules, whereby the simulated data is derived using historical maintenance data of the enterprise, historical sensor data collected by sensors in a factory of the enterprise. In this example, the digital twin interaction manager 60318 may receive requests for different maintenance schedules from a client device 60318 depicting an executive digital twin (e.g., a CFO digital twin, a CTO digital twin, or a CEO digital twin) and may initiate the simulations for each of the different maintenance schedules. The digital twin interaction manager 60318 may then serve the results of the simulation to the requesting client application.

[2300] In example embodiments, the digital twin interaction manager 60318 may manage one or more workflows that are performed via an executive digital twin. For example, the EMP 60000 may store a set of executive workflows, where each executive workflow corresponds to a role within an enterprise and includes one or more stages. In example embodiments, the digital twin interaction manager 60318 may receive a request to execute a workflow. The request may indicate the workflow and a user identifier. In response, the digital twin interaction manager 60318 may retrieve the requested workflow and may provide specific instructions and/or data to the client device 60102.

[2301] In example embodiments, the digital twin simulation system 60320 receives requests to run simulations using one or more digital twins. In example embodiments, the request may indicate a set of parameters that are to be varied and/or one or more simulation outcomes to output. In example embodiments, the digital twin simulation system 60320 may request one or more digital twins from the digital twin generation system 60310 and may vary a set of different parameters for

the simulation. In example embodiments, the digital twin simulation system 60320 may construct new digital twins and new data streams within existing digital twins. In example embodiments, the digital twin simulation system 60320 may perform environment simulation and/or data simulations. The environment simulation is focused on simulation of the digital twin ontology rather than the underlying data streams. In example embodiments, the digital twin simulation system 60320 generates simulated data streams appropriate for respective digital twin environments. This simulation allows for real world simulations of how a digital twin will respond to specific events such as changes in the cost of good supplied, or changes in the demand on the output of the factory.

[2302] In example embodiments, the digital twin simulation systems 60320 implement a set of models (e.g., physical mathematical forecasts, logical representations, or process diagrams) that develop the framework where data and the response of the digital twin can be simulated in response to different situational or contextual inputs/stimuli. In example embodiments, the digital twin simulation system 60320 may include or leverage a computerized model builder that constructs a predicted future state of either the data and/or the response of the digital twin to the input data. In example embodiments, the computerized model library may be obtained from a model data store 60338 that stores one or more models that define one or more behaviors of entities, such as based on scientific, economic, statistical, psychological, sociological, econometric, engineering, mathematical, physical, chemical, biological, architectural, computational, or other models, formulas, functions, processes, algorithms, or the like of the various types described herein or in the documents incorporated by reference herein (collectively referred to herein as "models" except where context indicates otherwise). The computerized digital twin model calculates the results of the model based on available inputs to build an interactive environment where users can watch and manipulate salient features of the simulated environment seeing how the entire system responds to specific changes in the environment. For example, the digital twin model may display how a set of objects that are stacked in a container will respond to tilting the container, where the behavior of the objects is based on a mechanical engineering or architectural model of the stacked objects, including structural features, weight distributions, and the like. This may assist in assessing the probability and/or impact of various fault modes, such as breaking, spilling, or the like, in response to seismic events, road conditions, weather conditions, wave action, or the like, as well as in simulating the response of other objects in the simulated environment, including in a chain of events. This may, for example, allow a user to identify events and consequences that occur as a result of multiple simultaneous or related faults or other events.

[2303] In example embodiments, digital twin behavior models may be updated and improved using results of actual experiments and real-world events. The use of such digital twin mathematical models and their simulations avoids actual experimentation, which can be costly and time-consuming. Instead, acquired knowledge about behavior of entities and computational power is used to diagnose and solve real-world problems cheaply and/or in a time-efficient manner. As such, the digital twin simulation system 60320 can facilitate understanding a system's behavior without actually testing the system in the real world. For example, to determine which type of wheel

configuration would improve traction the most while designing a tractor, a digital twin model simulation of the tractor could be used to estimate the effect of different wheel configurations on towing capacity. Useful insights about different decisions in the design may be gleaned without actually building the tractor. In addition, the digital twin simulation can support experimentation that occurs totally in software, or in human-in-the-loop environments where the digital twin represents systems or generates data needed to meet experiment objectives. Furthermore, digital twin simulations can be used to train persons using a perspective-appropriate virtual environment that would otherwise be difficult or expensive to produce.

[2304] In example embodiments, simulation environments may be constructed using a model capable of predicting a set of future states. These models include deep learning, regression models, quantum prediction engines, inference engines, pattern recognition engines, and many other forms of modeling engines that use historical outcomes, current state information, and other inputs to build a future state prediction. In example embodiments, a consideration in making the digital twin models' function is the ability to also show the response of the perspective-based digital twin structural elements (e.g., defining the deformation of the axle of a vehicle in response to different size loads). For example, the resultant digital twin representation can then be presented to the user in a virtual reality or augmented reality environment where specific perspectives are shown in their digital twin form.

[2305] In example embodiments, digital twins, as described herein, may be operating in coordination with an adaptive edge computing system and/or a set of adaptive edge computing systems that provide coordinated edge computation include a wide range of systems, such as classification systems (such as image classification systems, object type recognition systems, and others), video processing systems (such as video compression systems), signal processing systems (such as analog-to-digital transformation systems, digital-to-analog transformation systems, RF filtering systems, analog signal processing systems, multiplexing systems, statistical signal processing systems, signal filtering systems, natural language processing systems, sound processing systems, ultrasound processing systems, and many others), data processing systems (such as data filtering systems, data integration systems, data extraction systems, data loading systems, data transformation systems, point cloud processing systems, data normalization systems, data cleansing system, data deduplication systems, graph-based data storage systems, object-oriented data storage systems, and others), predictive systems (such as motion prediction systems, output prediction systems, activity prediction systems, fault prediction systems, failure prediction systems, accident prediction systems, event predictions systems, event prediction systems, and many others), configuration systems (such as protocol selection systems, storage configuration systems, peer-to-peer network configuration systems, power management systems, self-configuration systems, self-healing systems, handshake negotiation systems, and others), artificial intelligence systems (such as clustering systems, variation systems, machine learning systems, expert systems, rule-based systems, deep learning systems, and many others), system management and control systems (such as autonomous control systems, robotic control systems, RF spectrum management systems, network resource management systems, storage management

systems, data management systems, and others), robotic process automation systems, analytic and modeling systems (such as data visualization systems, clustering systems, similarity analysis systems, random forest systems, physical modeling systems, interaction modeling systems, simulation systems, and many others), entity discovery systems, security systems (such as cybersecurity systems, biometric systems, intrusion detection systems, firewall systems, and others), rules engine systems, workflow automation systems, opportunity discovery systems, testing and diagnostic systems, software image propagation systems, virtualization systems, digital twin systems, IoT monitoring systems, routing systems, switching systems, indoor location systems, geolocation systems, and others.

[2306] In example embodiments, the digital twin notification system 60322 provides notifications to users via enterprise digital twins associated with the respective users. In example embodiments, digital twin notifications are an important part of the overall interaction. The digital twin notification system 60322 may provide the digital twin notifications within the context of the digital twin setting so that the perspective view of the notification is set up specifically to enable enlightenment of how the notification fits into the general digital twin represented ontology, taxonomy, topology or the like.

[2307] FIG. 239 illustrates examples of different types of enterprise digital twins, including executive digital twins, in relation to the data layer, processing layer, and application layer of the enterprise digital twin framework. In example embodiments, executive digital twins may include, but are not limited to, CEO digital twins 60620, CFO digital twins 60622, COO digital twins 60624, CMO digital twins 60628, CTO digital twins 60630, CIO digital twins 60632, GC digital twins 60634, HR digital twins 60638, and the like. Additionally, the enterprise digital twins that may be relevant to the executive suite may include cohort digital twins 60640, agility digital twins 60642, CRM digital twins 60644, and the like. The discussion of the different types of digital twins is provided for example and not intended to limit the scope of the disclosure. It is understood that in example embodiments, users may alter the configuration of the various executive digital twins based on the business needs of the enterprise, the reporting structure of the enterprise, and the roles and responsibilities of the various executives within the enterprise.

[2308] In example embodiments, executive digital twins and the additional enterprise digital twins are generated using various types of data collected from different data sources. As discussed, the data may include real-time data 60660, historical data 60662, analytics data 60664, simulation/modeled data 60668, CRM data 60670, organizational data, such as org charts and/or an organizational digital twin 60672, an enterprise data lake 60674, and market data 60678. In example embodiments, the real-time data 60660 may include sensor data collected from one or more IoT sensor systems, which may be collected directly from each sensor and/or by various data collection devices associated with the enterprise, including readers (e.g., RFID, NFC, and BLUETOOTH™ readers), beacons, gateways, repeaters, mesh network nodes, WIFI systems, access points, routers, switches, gateways, local area network nodes, edge devices, and the like. Real-time data 60660 may include additional or alternative types of data that are collected in real-time, such as real-time sales data, real-time cost data, project management data that indicates the

status of current projects, and the like. Historical data may be any data collected by the enterprise and/or on behalf of the enterprise in the past. This may include sensor data collected from the sensor systems of the enterprise, sales data, cost data, maintenance data, purchase data, employee hiring data, employee on-boarding data, employee retention data, legal-related data indicating legal proceedings, patent filing data indicating patent filings and issued patents, project management data indicating historical progress of past and current projects, product data indicating products that are on the market, and the like. Analytics data 60664 may be data derived by performing one or more analytics processes on data collected by and/or on behalf of the enterprise. Simulation/modeled data 60668 may be any data derived from simulation and/or behavior modeling processes that are performed with respect to one or more digital twins. CRM data 60670 may include data obtained from a CRM of the enterprise. An organizational digital twin 60672 may be a digital twin of the enterprise. The enterprise data lake 60674 may be a data lake that includes data collected from any number of sources. Market data 60678 may be collected from many different sources and may be structured or unstructured. In example embodiments, market data 60678 may contain an element of uncertainty that may be depicted in a digital twin that relies on such market data 60678, such as by showing error bars, probability cones, random walk paths, or the like. It is appreciated that the different types of data highlighted above may overlap. For example, historical data may be obtained from the CRM data; the enterprise data lake 60672 may include real-time data 60660, historical data 60662, analytics data 60664, simulated/modeled data 460668, and/or CRM data 60670; and analytics data 60664 may be based on historical data 60662, real-time data 60660, CRM data 60670, and/or market data 60678. Additional or alternative types of data may be used to populate an enterprise digital twin.

[2309] In example embodiments, the data structuring system 60602 may structure the various data collected by and/or on behalf of the enterprise. In example embodiments, the digital twin generation system 60604 generates the enterprise digital twins. As discussed, the digital twin generation system 60604 may receive a request for a particular type of digital twin (e.g., a CEO digital twin 60620 or a CTO digital twin 60630) and may determine the types of data needed to populate the digital twin based on the configuration of the requested type of digital twin. In example embodiments, the digital twin generation system 60310 may then generate the requested digital twin based on the various types of data (which may include structured data structured by the data structuring system 60308). In example embodiments, the digital twin generation system 60310 may output the generated digital twin to a client application 60104, which may then display the requested digital twins.

[2310] In example embodiments, a CEO digital twin 60620 is a digital twin configured for the CEO or analogous top-level decision maker of an enterprise. The CEO digital twin 60620 may include high-level views of different states and/or operations data of the enterprise, including real-time and historical representations of major assets, processes, divisions, performance metrics, the condition of different business units of the enterprise, and any other mission-critical information type. In example embodiments, the CEO digital twin 60620 may work in connection with the EMP 60000 to provide simulations, predictions, statistical summaries, decision-support based on

analytics, machine learning, and/or other AI and learning-type processing of inputs (e.g., fiscal data, competitor data, product data, and the like). In example embodiments, a CEO digital twin 60620 may provide functionality including, but not limited to, management of personnel, delegation of tasks, decisions, or tasks, coordination with the Board of Directors and/or strategic partners, risk management, policy management, oversight of budgets, resource allocation, investments, and other executive-related resources.

[2311] In example embodiments, the types of data that may populate a CEO digital twin 60620 may include, but are not limited to: macroeconomic data, microeconomic analytic data, forecast data, demand planning data, employment and salary data, analytic results of AI and/or machine learning modeling (e.g., financial forecasting), prediction data, recommendation data, securities-relevant financial data (e.g., earnings, profitability), industry analyst data (e.g., Gartner quadrant), strategic competitive data (e.g., news and events regarding industry trends and competitors), business performance metrics by business unit that may be relevant to evaluating performance of the business units (e.g., P&L, head count, factory health, supply chain metrics, sales metrics, R&D metrics, marketing metrics, and many others), Board package data, or some other type of data relevant to the operations of the CEO and/or executive department. In example embodiments, the digital twin system 60004 may obtain securities-relevant financial data from, for example, the enterprise's accounting software (e.g., via an API), publicly disclosed financial statements, third-party reports, tax filings, and the like. In example embodiments, the digital twin system 60004 may obtain strategic competitive data from public news sources, from publicly disclosed financial reports, and the like. In example embodiments, macroeconomic data may be derived analytically from various financial and operational data collected by the EMP 60000. In example embodiments, the business performance metrics may be derived analytically, based at least in part on real time operations data, by the artificial intelligence services system 60012 and/or provided from other users and/or their respective executive digital twins. The CEO digital twin 60620 may be used to define real time operations data parameters of interest and to monitor, collect, analyze, and interpret real time operations data for conformance to and alignment with an organization's stated business objects, Board requirements, industry best practice, regulation, or some other criterion.

[2312] In example embodiments, a CEO digital twin 60620 may include high-level views of different states of the enterprise, including real-time and historical representations of major assets, the condition of different business units of the enterprise, and any mission-critical information. The CEO digital twin 60620 may initially depict the various states at a lower granularity level. In example embodiments, a user that is viewing the CEO digital twin 60620 may select a state to drill down into the selected state and view the selected state at a higher level of granularity. For example, the CEO digital twin 60620 may initially depict a subset of the various states of the enterprise at a lower granularity level, including a financial-department state (e.g., a visual indicator indicating an overall financial health score of the enterprise). In response to selection, the CEO digital twin 60620 may provide data, analytics, summary, and/or reporting including, but not limited to, real-time, historical, aggregated, comparison, and/or forecasted financial information (e.g., real-time, historical, simulated, and/or forecasted revenues, liabilities, and the like). n this way, the CEO

digital twin 60620 may initially present the user (e.g., the CEO) with a view of various different aspects of the enterprise (e.g., different indicators to indicate different "health" levels of a respective business unit or part of the enterprise) but may allow the user to select which aspects require more of her attention. In response to such a selection, the CEO digital twin 60620 may request a more granular view of the selected state(s) from the EMP 60000, which may return the requested states at the more granular level.

[2313] In example embodiments, a CEO digital twin 60620 may include an executive-level digital twin of the executive department (e.g., C-suite, directors, Board members, and the like), which the user may use to identify, assign, instruct, oversee and review executive department personnel and third-party personnel, departments, organizations and the like that are associated with the activities of the executive of an organization, including the Board of Directors and the like that are involved in the oversight of the organization's management. In example embodiments, the executive-level digital twin may include a definition of the various roles, employees, and departments working under the CEO, the reporting structure for each individual in the business unit and may be populated with the various names and/or other identifiers of the individuals filling the respective roles. In example embodiments, the CEO digital twin 60620 may include a graphical user interface that provides the user the ability to define/redefine personnel groupings, assign performance criteria and metrics to business units, roles, and/or individuals, and/or assign/delegate tasks to business units, roles, and/or individuals, and the like via the executive-level digital twin. In example embodiments, the executive-level digital twin may provide real-time operations data of the organization to continuously evaluate the personnel groupings' performance against the stored performance criteria.

[2314] In example embodiments, a CEO digital twin 60620 may be configured to interface with the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the executive department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools (e.g., where the collaboration occurs to some extent within a common interface by which the digital twin entities are viewed and collaboration activities take place and/or where the components of the EMP that used to configure, operate or support the digital twin also govern collaboration around digital twin entities and workflows), whiteboard tools, agile development environment tools (such as features in Slack™ environments), presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein. The collaboration tools may include collaborative communication (e.g., facilitating live conferencing where participants are simultaneously presented with conference-related views of digital twin entities or workflows), asynchronous collaboration (such as where actions on digital twin entities, comments, or the like are represented to different users who interact with the entities), version control features, and many others.

[2315] In example embodiments, a CEO digital twin 60620 may be configured to provide research, track, and report on an executive department initiative including, but not limited to, an overall

strategic goal, policy implementation, product roll-out, Board interaction, investment or acquisition, investor relations, public relations and press handling, budgeting, or some other type of executive initiative. The CEO digital twin 60620 may interact with and share such data and reporting with other executive digital twins, including, but not limited to, a CFO digital twin, a COO digital twin, and the like. In example embodiments, the CEO digital twin 60620 or an executive agent integrated with or within it (such as one trained to undertake expert executive actions as described elsewhere herein) may leverage intelligence services (e.g., data analytics, machine learning and A.I. processes) to analyze financial reports, projections, simulations, budgets, and related summaries to identify key departments, personnel, third-party or others that are, for example, listed in, or subject to, a project, initiative, budget line item and the like, and who therefore may have an interest in such material. Such material pertaining to a given party may be abstracted and summarized for presentation, and formatted and presented automatically, or at the direction of the CEO or other user, to the party that is the origin of the expense and/or subject of the material. For example, the CEO digital twin 60620 may assemble materials for the purposes of developing presentations, speaking points, press releases, or some other material for the CEO or other executive personnel to use for public presentation. In examples, a CEO in anticipation of giving a conference presentation on the introduction of a new company product may use the CEO digital twin 60620 to specify and configure the identification, collection and assembly of operations data that is relevant to the upcoming presentation, such as product data (e.g., units produced, units shipped), financial data (e.g., products sold, products reserved), graphic presentation information (e.g., product photos, maps of product distribution, graphs of anticipated sales), forecasting data (e.g., market growth expected), or some other type of data and assemble such information in a presentation format, such as presentation slides, white paper template, speech talking points, press release, or some other summary format that may form the basis of the presentation or be distributed in conjunction with the presentation and/or its marketing.

[2316] In example embodiments, a CEO digital twin 60620 may be configured to track and report on stakeholder communications (e.g., reports, Board requests, investor requests) related to the executive department. The CEO digital twin 60620 may present, store, analyze, reconcile and/or report on executive activities related to parties with whom the executive department is contracting, cooperating with, reporting to and so forth, such as key personnel, outside contractors, the press, the Board of Directors, or others.

[2317] In example embodiments, the CEO digital twin 60620 may be configured to simulate one or more aspects of the enterprise. Such simulations may assist the user (e.g., the CEO) in making executive level decisions. Simulations of a proposed executive initiative may be tested, for example using the modeling, machine learning, and/or AI techniques, as described herein, by simulating temporal effects on initiatives (e.g., introduction of a new product), varying financial parameters (e.g., potential investment levels), targeting parameters (e.g., geographic, demographic, or the like), and/or other suitable executive parameters. In example embodiments, the digital twin simulation system 60320 may receive a request to perform an executive simulation requested by the CEO digital twin 60620, where the request indicates one or more parameters that are to be

varied in one or more enterprise digital twins. In response, the digital twin simulation system 60320 may return the simulation results to the CEO digital twin 60620, which in turn outputs the results to the user via the client device display. In this way, the user may be provided with various outcomes corresponding to different parameter configurations. In example embodiments, an executive agent may be trained to recommend and/or select a parameter set based on the respective outcomes associated with each respective parameter set.

[2318] In example embodiments, a CEO digital twin 60620 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to an executive strategy, executive planning, executive activities, and/or executive initiatives. For example, the CEO digital twin 60620 may be associated with a plurality of databases or other repositories of financial materials, summaries and reports and analytics, including such materials, summaries and reports and analytics related to prior executive activity (e.g., prior quarterly financial performance, prior investments, prior strategic partners, co-developments, and the like), each of which may be further associated with financial and performance metrics pertaining to the campaign and which are also accessible to the CEO digital twin 60620.

[2319] In example embodiments, a CEO digital twin 60620 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to financial reporting, ratings, rankings, financial trend data, income data, or other data related to an executive's responsibilities. A CEO digital twin 60620 may link to, interact with, and be associated with external data sources, and able to upload, download, aggregate external data sources, including with the EMP's internal data, and analyze such data, as described herein. Data analysis, machine learning, AI processing, and other analysis may be coordinated between the CEO digital twin 60620 and an analytics team based at least in part on using the artificial intelligence services system 60012. This cooperation and interaction may include assisting with seeding executive-related data elements and domains in the enterprise data store 60014 for use in modeling, machine learning, and AI processing to identify an optimal business strategy, or some other executive-relating metric or aspect, as well as identification of the optimal data measurement parameters on which to base judgment of an executive initiative's success. Examples of data sources 60030 that may be connected to, associated with, and/or accessed from the CEO digital twin 60620 may include, but are not limited to, a sensor system 60032 having sensors that sensor data from facilities (e.g., manufacturing facilities, shipping and logistics facilities, transportation facilities, agricultural facilities, resource extraction facilities, computing facilities, and many others) and/or other physical entities of the enterprise, a sales database 60034 that is updated with sales figures in real time, a CRM system 60038, a content marketing platform 60040, an executive resource planning system 60034, financial databases 60048, surveys 60050, org charts 60052, workflow management systems 60054, third-party data sources 60060, customer databases 60062 that store customer data, and/or third-party datastores 60060 that store third-party data, edge devices 60064 that report data relating to physical assets (e.g., smart machinery/manufacturing equipment, sensor kits, autonomous vehicles of the enterprise, wearable devices, and the like), enterprise management systems 60080, HR systems 60082, content management systems 60084, and the like). In example embodiments, the digital

twin system 60004 abstracts the different views (or states) within the digital twin to the appropriate granularity. For instance, the digital twin system 60004 may have access to all the sensor data collected on behalf of the enterprise as well as access to real-time sensor data streams. Typically, such data is far too granular for an executive such as a CEO, and sensor data readings are often of little importance to the CEO unless associated with a mission critical state or operation. In this example, however, if the sensor readings from a particular physical asset (e.g., a critical piece of manufacturing equipment) are indicative of a potentially critical situation (e.g., failure state, dangerous condition, or the like), then the analytics that indicate the potentially critical situation may become very important to the CEO. Thus, the digital twin system 60004 may, when building the appropriate perspective for the CEO, include a state indicator of the physical asset in the CEO digital twin 60620. In this way, the CEO can drill down into the state indicator of the physical asset to view the potentially critical situation at a greater granularity (e.g., the machinery and an analysis of the sensor data used to identify the situation).

[2320] In example embodiments, a CEO digital twin 60620 may be configured to monitor an organization's performance based at least in part on real time operations data and the use of the monitoring agent of the client application 60104, as described herein, that is associated with the CEO digital twin 60620. The monitoring agent may report on such activities to the EMP 60000 for presentation in a user interface that is associated with the CEO digital twin 60620. In response, the EMP 60000 may train an executive agent (which may include one or more machine-learned models) to handle and process such notifications when they next arrive, and escalate and/or alert the CEO when such notifications are of an urgent nature, such as an announcement of an acquisition by a competitor, a report indicating an under-performing business unit, a high-profile press article, a radical change in the stock market (for the CEO's company, a cohort member, or the market as a whole), a downgrade in rating by an industry analyst, an external event likely to disrupt operations (such as a natural disaster or epidemic) or some other important event. In example embodiments, the CEO digital twin 60620 may generate performance alerts based on real time operations data, performance trends, and the like. This may allow a CEO to optimize initiatives in real-time without having to manually request such real-time data; the CEO digital twin 60620 may automatically present such information and related/necessary alerts as configured by the organization, CEO, or some other interested party.

[2321] In example embodiments, a CEO digital twin 60620 may be configured to report on the performance of the executive department, personnel of the executive department, executive activities, executive content, executive platforms, executive partners, or some other aspect of management within a CEO's responsibilities. Reporting may be to the CEO, the executive department, to other executives of an organization (e.g., the COO), or to outside third parties (e.g., partners, press releases, and the like). As described herein, reporting may include stakeholder summaries, minutes of meetings, presentations, sales data, customer data, financial performance metrics, personnel metrics, data regarding resource usage, industry summaries (e.g., summaries of merger and acquisition activity in an industry segment), or some other type of reporting data. Reporting and the content of reporting may be shared by the CEO digital twin 60620 with other

executive digital twins. The reporting functionality of the CEO digital twin 60620 may also be used for populating new or preset reporting formats, and the like. Templets of common reporting formats may be stored and associated with the CEO digital twin 60620 to automate the presentation of data and analytics according to pre-defined formats, styles and system requirements. In example embodiments, an executive agent trained by the user may be trained to surface the most important reports to the user. For example, if the user (e.g., the CEO) consistently views and follows up on sales data reports but routinely skips over reports relating to the manufacturing KPIs, the executive agent may automatically surface sales data reports to the user and may automatically delegate manufacturing KPIs to another executive digital twin (e.g., the COO digital twin).

[2322] In example embodiments, a CEO digital twin 60620 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to competitors of a CEO's organization, or named entities of interest. In example embodiments, such data may be collected by the EMP 60000 via data aggregation, spidering, web-scraping, or other techniques to search and collect competitor information from sources including, but not limited to, information on investment and/or acquisitions, press releases, SEC or other financial reports, or some other publicly available data. For example, a user wishing to monitor a certain competitor may request that the CEO digital twin 60620 provide materials relating to the certain competitor. In response, the EMP 60000 may identify a set of data sources that are either publicly available or to which the enterprise of the CEO has access (e.g., internal data sources, licensed third-party data, or the like). The EMP 60000 may configure a cohort digital twin 60640 based on the types of data/analysis/services the user requests and the identified set of data sources. The EMP 60000 may then serve the cohort digital twin 60640 associated with the requested party (e.g., competitor) to the CEO digital twin 60620.

[2323] In example embodiments, a CEO digital twin 60620 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to regulatory activity, such as government regulations, industry best practices or some other requirements or standards. For example, the CEO digital twin 60620 may be in communication with another enterprise digital twin, such as a General Counsel digital twin 60634, through which the legal team can keep the CEO apprised of new regulation or regulation changes as they occur.

[2324] In example embodiments, the client application 60104 that executes the CEO digital twin 60620 may be configured with an executive agent 60704 that is trained on the CEO's actions (which may be indicative of behaviors, and/or preferences). In example embodiments, the executive agent 60704 may record the features relating to the actions (e.g., the circumstances relating to the user's action) to the expert agent system 60010. For example, the executive agent 60704 may record each time the user delegates a task to a subordinate (which is the action) as well as the features surrounding the delegation of the task (e.g., an event that caused the user to delegate the task, the type of task that was delegated, the role to which the task was delegated, instructions provided by the user with the delegation, and the like). The executive agent 60704 may report the actions and features to the expert agent system 60010 and the expert agent system 60010 may train the executive agent 60704 on the manner by which the executive agent 60704 can delegate or

recommend delegation of tasks in the future. Once trained, the executive agent 60704 may automatically perform actions and/or recommend actions to the user. Furthermore, in example embodiments, the executive agent 60704 may record outcomes related to the performed/recommended actions, thereby creating a feedback loop with the expert agent system 60010.

[2325] References to features and functions of the EMP and digital twins in this example of a CEO digital twin 60620 should be understood to apply to other digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2326] In example embodiments, a Chief Financial officer (CFO) digital twin 60622 may be a digital twin configured for a CFO of an enterprise, or an analogous executive tasked with overseeing the finance-related tasks of the enterprise. A CFO digital twin 60622 may provide data, analytics, summary, and/or reporting including, but not limited to, real-time, historical, aggregated, comparison, and/or forecasted financial information (e.g., real-time, historical, simulated, and/or forecasted sales figures, expenditures, revenues, liabilities, and the like). In example embodiments, the CFO digital twin may work in connection with the EMP 60000 to provide simulations, predictions, statistical summaries, decision support based on analytics, machine learning, and/or other AI and learning-type processing of inputs (e.g., accounting data, sales data, sensor data and the like).

[2327] In example embodiments, a CFO digital twin 60622 may provide features and functionality including, but not limited to, management of financial personnel, partners and outside consultants and contractors (e.g., accounting firms, auditors and the like), oversight of budgets, procurement, expenditures, receivables, and other finance-related resources, compliance, oversight of sales and sales staff and departments' financial performance, management of contracting, management of internal policies (e.g., policies related to expenditures and reporting), tax law, finance-related privacy law (e.g., pertaining to credit agency data), reporting, compliance, and regulatory analysis.

[2328] In example embodiments, the types of data that may populate a CFO digital twin may include, but are not limited to, financial performance metrics by business unit, by product, by geography, by factory, by store location(s), by asset class, earnings, cash, balance sheet data, cash flow, profitability, resource utilization, audit data, general ledger data, asset performance data, securities and commodities data, insurance and risk management data, asset aging and depreciation data, asset allocation data, macroeconomic data, microeconomic analytic data, tax data, pricing data, competitive product and pricing data, forecast data, demand planning data, employment and salary data, analytic results of AI and/or machine learning modeling (e.g., financial forecasting), prediction data, recommendation data, or some other type of data relevant to the operations of the CFO and/or finance department. In example embodiments, "datum," "data," "dataset," "datastore," "data warehouse," and/or "database," as used herein, may refer to information that is stored in a numeric or statistical format, including summaries, inputs or outputs in statistical or scientific notation, and also includes information that is stored in natural language format (e.g., text excerpts from reports, press releases, statutes and the like), information in a graphic format (e.g., financial performance graphs), information in audio and/or audio-visual format (e.g., recorded audio from

conference calls or video from presentations, including natural language transcript summaries of audio and/or audio-visual formatted information), or some other type of information.

[2329] In example embodiments, a CFO digital twin 60622 may depict a finance department twin of the finance department, which the user may use to identify, assign, instruct, oversee and review finance department personnel and third-party personnel that are associated with the finance activities of an organization, including third-party partners and other outside contractors, such as accounting firms, tax lawyers and the like that are involved in the organization's finance endeavors. Examples of such organization personnel include, but are not limited to, finance department staff, sales analysts, statisticians, data scientists, executive personnel, human resources staff, Board Members, advisors, or some other type of organization personnel relevant to the functioning of a finance department. Examples of a finance department's third-party personnel include, but are not limited to, lawyers, accountants, management consultants, social media platform personnel, finance partners, consultants, contractors, financial firm staff, auditors, or some other type of third-party personnel.

[2330] In example embodiments, the CFO digital twin 60622 may include a definition of the various roles/employees working under the CFO, the reporting structure, and associated permissions, for each individual in the business unit, and may be populated with the various names and/or other identifiers of the individuals filling the respective roles. In example embodiments, a user (e.g., the CFO of an enterprise) may use the CFO digital twin 60622 to adjust the reporting structure within the finance department and/or to grant permissions to one or more individuals within the department.

[2331] In example embodiments, a CFO digital twin 60622 may be configured to interface with the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the finance department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein.

[2332] In example embodiments, a CFO digital twin 60622 may be configured to research, create, track and report on a finance department initiative including, but not limited to, an overall department budget, a budget for a single or group of finance initiatives, an audit, a third-party vendor activity, or some other type of expense or budget. In example embodiments, the CFO digital twin 60622 may interact with and share such expense or budget data and reporting with other enterprise twins, as described herein, including, but not limited to, a digital twin related to accounts payable, executive staff such as the CEO (e.g., CEO digital twin) or COO (e.g., the COO digital twin), or other suitable enterprise digital twins. In example embodiments, the CFO digital twin 60622 may leverage one or more intelligence services of the EMP 60000 based at least in part on the data analytics, machine learning and A.I. processes, as described herein, to provide financial reports, projections, simulations, budgets and related summaries. In some of these embodiments, the CFO digital twin 60622 my use the intelligence services to identify key departments, personnel,

third-party or others that are, for example, listed in, or subject to, the budget line item and who therefore may have an interest in such material. Budget material pertaining to a given party may be abstracted and summarized for presentation independent from the entirety of the budget, and formatted and presented automatically, or at the direction of the CFO or other user, to the party that is the origin of the expense and/or subject of the budget item.

[2333] In example embodiments, a CFO digital twin 60622 may be configured to track and report on inbound and outbound billing (i.e., accounts receivable and payable) related to the finance department and/or organization. In example embodiments, the CFO digital twin 60622 may include a billing digital twin that identifies the billing department, personnel, processes and systems associated with the billing workflows of the enterprise. In these embodiments, the billing digital twin may interact present, store, analyze, reconcile and/or report on billing activities related to parties with whom the finance department is interacting. In example embodiments, the user of the CFO digital twin 60622 may approve bills, issue bills, drill down into a set of bills, initiate investigations of bills or the like via the GUI if the CFO digital twin 60622.

[2334] In example embodiments, a CFO digital twin 60622 may be configured to provide a user (e.g., a CFO or other finance department executive) with information that is unique to the CFO digital twin 60622 and thus can provide insights and perspectives on financial performance that are unique to the CFO digital twin 60622. In example embodiments, the EMP 60000 in connection with the CFO digital twin 60622 may create and derive new financial metrics and analytics including, but not limited to, functionalities such as native data and model creation, and data and model combinations and aggregations based at least in part on the real time operations of an organization. Native data and model creation, such as specifying the data to be collected, the format within which to collect and store the data, the data transformations to model, and so forth gives one the ability to craft, combine, aggregate, modify, transform, and/or weight the native data (including in combination with other third-party data) in manners that are appropriately mathematically tuned to the modeling, analysis, machine learning, and/or AI techniques that are performed by the EMP 60000 and CFO digital twin 60622, rather than being reliant on data and/or model presets. Similarly, in the analytic context of the CFO's operations and the function of the EMP and CFO digital twin 60622, native data and model creation and structuring by the EMP and CFO digital twin 60622 enables analytics, machine learning, AI operations and the like, yielding new analytic results and insights, based at least in part on the real time operations of an organization, because the EMP and CFO digital twin 60622 has enabled the CFO to move further up in financial data creation and modeling operations to assert greater creative control over the types of data and other input material to be used in developing analytic insights that may be created and reported for the purpose of improving performance including, but not limited to, product margins (e.g., gross, contribution, net and the like), product features, upsell opportunities or some other performance metric.

[2335] In example embodiments, the CFO digital twin 60622 may be configured to simulate finance-related activities on behalf of a user. In these embodiments, the user may identify one or more parameters that can be varied during for a simulation including, but not limited to, financial

and/or budget parameters, pricing and sales goal settings, process designs, and maintenance/infrastructure upgrades, internal controls design, product testing frequencies/types, manufacturing down-times, flexible workforce planning, and the like. In these embodiments, the digital twin simulation system 60320 may receive a request to perform the simulation requested by the CFO digital twin 60622, where the request indicates features and the parameters, including financial parameters, that are to be varied. In response, the digital twin simulation system 60320 may return the simulation results to the CFO digital twin 60622, which in turn outputs the results to the user via the client device display. In this way, the user is provided with various outcomes corresponding to different parameter configurations. In example embodiments, the user may select a parameter set based on the various outcomes. In example embodiments, an executive agent trained by the user may select the parameter sets based on the various outcomes. The simulations, analytics and/or modeling performed by the CFO digital twin 60622 may be used to mitigate risk for IPO, M&A, equity and debt offerings, or some other type of transaction. The simulations, analytics and/or modeling performed by the CFO digital twin 60622 may be used to create and structure sales incentives, including commissions and other performance-based compensation. The simulations, analytics and/or modeling performed by the CFO digital twin 60622 may be used to evaluate insurance offerings and other information related to business interruption preparedness. The simulations, analytics and/or modeling performed by the CFO digital twin 60622 may be used to analyze loan covenant monitoring and projections. The CFO equipped with digital twin 60622 will be better able to adapt quickly to change by predicting headwinds, forecasting operational performance, and making informed decisions across departments while mitigating risk.

[2336] In example embodiments, a CFO digital twin 60622 may be configured to access insights across environmental resource management (ERM) solutions for risk oversight that includes, but is not limited to, internal controls design, testing, certification, and reporting while directing listed actions into a repository. In example embodiments, a CFO digital twin 60622 may be configured to streamline governance, risk management, and compliance processes in order to connect risk and compliance across the organization and manage complex audit fieldwork and work papers.

[2337] In example embodiments, a CFO digital twin 60622 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to a financial strategy, plan, activity or initiative. For example, the CFO digital twin 60622 may be associated with a plurality of databases or other repositories of financial materials, summaries and reports and analytics, including such materials, summaries and reports and analytics related to prior financial activity (e.g., prior quarterly financial performance), each of which may be further associated with third-party financial or economic data.

[2338] In example embodiments, a CFO digital twin 60622 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to financial reporting, ratings, rankings, financial trend data, income data, or other finance department-related data. A CFO digital twin 60622 may link to, interact with, and be associated with external data sources, and able to upload, download, aggregate external data sources, including with the EMP's internal data, and analyze such data. Data analytics, machine learning, AI processing, and other data-driven

processes may be coordinated between the CFO digital twin 60622 and an analytics team based at least in part on insights derived by the artificial intelligence services system 60012. This cooperation and interaction may include assisting with seeding finance-related data elements and domains in the enterprise data store 60014 for use in modeling, machine learning, and AI processing to identify the optimal financial strategy, or some other finance-related metric or aspect, as well as identification of the optimal data measurement parameters on which to base judgment of a finance endeavor's success. Examples of data sources 60030 that may be connected to, associated with, and/or accessed from the CFO digital twin 60622 may include, but are not limited to, a sensor system 60032, a sales database 60034 that is updated with sales figures in real time, a CRM system 60038, a finance platform 60040, news websites, a financial database 60048 that tracks costs of the business, org charts 60052 , a workflow management system 60054, customer databases 60062 that store customer data, and/or third-party datastores 60060 that store third-party data.

[2339] In example embodiments, a CFO digital twin 60622 may aggregate data sources and types, creating new data types, summaries and reports that are not available elsewhere. This may reduce reliance upon the need of multiple third-party providers and current solutions. This may, among other benefits and improvements, reduce expenses associated with acquiring data needed for sound financial decision making.

[2340] In example embodiments, a CFO digital twin 60622 may be configured to monitor a user's performance of finance-related tasks via a monitoring function of an agent of the client application 60104 executing the CFO digital twin 60622. In example embodiments, the monitoring function of the executive agent may report on certain activities to the EMP 60000 that are undertaken by the user when interfacing with the CFO digital twin 60622. In response, the EMP 60000 may train the executive agent (which may include one or more machine-learned models) to handle and process such finance-related tasks when they next arrive. For example, the monitoring function may monitor when the user (e.g., the CFO) escalates a state of the CFO digital twin 60622 to the CEO and/or when the user delegates a task to a subordinate via the CFO digital twin 60622. Each time such escalations and/or delegation events occur and/or when the user (e.g., the CFO or other finance executive) responds to an alert or other notifications of an urgent nature and may report the actions taken by the user in response to each respective account to the EMP 60000. In response, the executive agent system 60010 may train an executive agent 60704 based on the reported actions, which in turn may be leveraged by the CFO digital twin 60622 to respond to certain later occurring events on which the executive agent 60704 was trained on (e.g., analytics showing poor financial performance or finance activity (e.g., a new investment). For example, an executive agent 60704 trained with respect to a CFO digital twin 60622 may automatically issue financial performance alerts to certain employees based on performance trends of one or more business units. In another example, the executive agent 60704 may automatically escalate a notification to the CEO (which may be depicted in the CEO digital twin 60620) when certain metrics indicate a poor financial forecast. In example embodiments, the executive agent 60704 in connection with the CFO digital twin 60622 may allow a CFO to optimize initiatives in real-time without having

to manually request such real-time financial performance data. In example embodiments, the CFO digital twin 60622 may automatically present such information and related/necessary alerts as configured by the configuring user, the CFO, or some other user having such permissions.

[2341] In example embodiments, an executive agent 60704 trained in connection with a CFO digital twin 60622 may be configured to report on the performance of the finance department, personnel of the finance department, finance activities, finance content, finance platforms, finance partners, or some other aspect of management within a CFO's responsibilities. Reporting may be to the CEO, the Board of Directors, other executives of an organization (e.g., the COO), or to outside third parties (e.g., partners, press releases, and the like). The reporting functionality of the CFO digital twin 60622 may also be used for populating required data for formal reporting requirements such as shareholder statements, annual reports, SEC filings, and the like. Templets of common reporting formats may be stored and associated with the CFO digital twin 60622 to automate the presentation of data and analytics according to pre-defined formats, styles and system requirements.

[2342] In example embodiments, a CFO digital twin 60622 in combination with the EMP 60000 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to competitors of a CFO's organization, or named entities of interest. In example embodiments, such data may be collected by the EMP 60000 via data aggregation, spidering, web-scraping, or other techniques to search and collect competitor information from sources including, but not limited to, press releases, SEC or other financial reports, mergers and acquisitions activity, or some other publicly available data.

[2343] In example embodiments, a CFO digital twin 60622 in combination with the EMP 60000 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to regulatory activity, such as government regulations, industry best practices or some other requirements or standards. For example, the CFO digital twin 60622 may be in communication with another enterprise digital twin, such as a General Counsel digital twin 60634, through which the legal team can keep the CFO apprised of new regulations or regulation changes as they occur.

[2344] In example embodiments, the client application 60104 that executes the CFO digital twin 60622 may be configured with an executive agent that reports a CFO's behaviors and preferences (or other finance personnel's behaviors and preferences) to the expert agent system 60010, as described herein, and the expert agent system 60010 may train the executive agent on how the CFO or other finance personnel respond to certain situations and adjust its operation based at least in part on the data collection, analysis, machine learning and A.I. techniques, as described herein. The foregoing examples are optional examples and are not intended to limit the scope of the disclosure.

[2345] References to features and functions of the EMP and digital twins in this example of a finance department and a CFO digital twin 60622 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2346] In example embodiments, a Chief Operating Officer (COO) digital twin 60624 may be a digital twin configured for a COO of an enterprise, or an analogous executive tasked with overseeing the operations tasks of the enterprise. A COO digital twin 60624 may provide functionality including, but not limited to, management of personnel and partners, oversight of various departments (e.g., oversight over marketing department, HR department, sales department, and the like), project management, implementation and/or rollouts of business processes and workflows, budgeting, reporting, and many other operations-related tasks.

[2347] In example embodiments, a COO digital twin 60624 may provide data, analytics, summary, and/or reporting including, but not limited to, real-time, historical, aggregated, comparison, and/or forecasted financial information (e.g., sales, expenditures, revenues, liabilities, profitability, cash flow and the like), mergers and acquisitions information, systems data, reporting and controls data, or some other operations related information. In example embodiments, the COO digital twin 60624 may work in connection with the EMP 60000 to provide simulations, predictions, statistical summaries, decision support based on analytics, machine learning, and/or other AI and learning-type processing of inputs (e.g., equipment data, sensor data and the like), for example, those related to the development, communication and implementation of effective growth strategies and processes for an organization.

[2348] In example embodiments, the types of data that may populate a COO digital twin may include, but are not limited to, operations data, key performance indicators (KPIs) for factories/plants, business units, assets/equipment; uptime/downtime, safety data, risk management data, demand plan data, logistics data, workflow data, financial performance metrics by business unit, by product, by geography, by factory, by store location(s), by asset class, earnings, resource utilization; audit data, asset performance data, asset aging and depreciation data, asset allocation data, or some other type of operations-relevant data or information.

[2349] In example embodiments, a COO digital twin 60624 may depict a twin of the operations department, which the user may use to identify, assign, instruct, oversee and review operations department personnel and third-party personnel that are associated with the design, implementation and evaluation of operational processes, internal infrastructures, reporting systems, company policies, and the like.

[2350] In example embodiments, the COO digital twin 60624 may include a definition of the various roles/employees working under the COO, the reporting structure, and associated permissions, for each individual in the business unit, and may be populated with the various names and/or other identifiers of the individuals filling the respective roles.

[2351] In example embodiments, a COO digital twin 60624 may be configured to interface with the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the operations department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein.

[2352] In some of these embodiments, the COO digital twin 60624 may be configured to simulate operations activities, such as a proposed new operational plan, process or program. In these embodiments, the digital twin simulation system 60320 may receive a request to perform the simulation requested by the COO digital twin 60624, where the request indicates features and the parameters of the operational plan or other activity that is proposed for implementation, the associated variables for which may be altered or varied to produce differing simulation environments. In response, the digital twin simulation system 60320 may return the simulation results to the COO digital twin 60624, which in turn outputs the results to the user via the client device display. In this way, the user is provided with various outcomes corresponding to different operational parameter configurations. In example embodiments, an executive agent trained by the user may select the parameter sets based on the various outcomes.

[2353] In example embodiments, a COO digital twin 60624 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to an operations strategy, plan, activity or initiative. For example, the COO digital twin 60624 may be associated with a plurality of databases or other repositories of operational data, summaries and reports and analytics, including such materials, summaries and reports and analytics related to prior operations activity, each of which may be further associated with financial and performance metrics pertaining to the activity and which are also accessible to the COO digital twin 60624.

[2354] In example embodiments, a COO digital twin 60624 may be configured to monitor operational performance, including in real time, based at least in part on use of the monitoring agent of the client application 60104, as described herein, that is associated with the COO digital twin 60624. The monitoring agent may report on such activities to the EMP 60000 for presentation in a user interface that is associated with the COO digital twin 60624. In response, the EMP 60000 may train an executive agent (which may include one or more machine-learned models) to handle and process such notifications when they next arrive and escalate and/or alert the COO when such notifications are of an urgent nature.

[2355] In example embodiments, a COO digital twin 60624 may be configured to report on the performance of the operations department, personnel of the operations department, operations activities, operations content, operations platforms, operations partners, or some other aspect of management within a COO's responsibilities.

[2356] In example embodiments, the EMP 100 trains and deploys executive agents on behalf of enterprise users. In example embodiments, an executive agent is an AI-based software system that performs tasks on behalf of and/or suggests actions to a respective executive user. In example embodiments, the EMP 100 receives data from various data sources associated with a particular entity or workflow and learns the workflows performed by the particular user based on the data and the surrounding circumstances or context. For example, the user may be a COO that is presented with a COO digital twin 60624. Among the responsibilities of the COO may be scheduling maintenance and replacement of equipment in a manufacturing, warehouse, or other operational facilities. The states depicted in the COO digital twin 60624 may include depictions of the condition of different pieces of equipment within the operational factory. In this example, the

COO may schedule maintenance via the digital twin when a piece of equipment is determined to be in a first condition (e.g., a deteriorating condition) and may issue a request to the COO via the COO digital twin 60624 to replace the piece of equipment when the equipment is determined to be in a second condition (e.g., a critical condition). The executive agent may learn the COO's tendencies based on the COO's previous interaction with the COO digital twin 60624. Once trained, the executive agent may automatically request replacements from the COO when a particular piece of equipment is determined to be in the second condition and may automatically schedule maintenance if the piece of equipment is in the first condition.

[2357] In example embodiments, the client application 60104 that executes the COO digital twin 60624 may be configured with an executive agent that reports a COO's behaviors and preferences (or other operations personnel's behaviors and preferences) to the executive agent system 60010, as described herein, and the executive agent system 60010 may train the executive agent on how the COO or other executive personnel respond to certain situations and adjust its operation based at least in part on the data collection, analysis, machine learning and A.I. techniques, as described herein. The foregoing examples are optional examples and are not intended to limit the scope of the disclosure.

[2358] References to features and functions of the EMP and digital twins in this example of an operations department and a COO digital twin 60624 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2359] In example embodiments, a Chief Marketing Officer (CMO) digital twin 60628 may be a digital twin configured for a CMO of an enterprise, or an analogous executive tasked with overseeing the marketing tasks of the enterprise. A CMO digital twin 60628 may provide functionality including, but not limited to, management of personnel and partners, development and oversight of marketing budgets and resources, management of marketing and advertising platforms, development and management of marketing content, strategies and campaigns, reporting, competitor analysis, regulatory analysis, and management of data privacy and security.

[2360] In example embodiments, the types of data that may populate and/or be utilized by a CMO digital twin 60628 may include, but are not limited to, macroeconomic data; market pricing data; competitive product and pricing data; microeconomic analytic data; forecast data; demand planning data; competitive matrix data; product roadmap; product capability data; consumer; consumer profile data; collaborative filtering data; analytic results of AI and/or machine learning modeling; channel data; demographic data; geographic data; prediction data; recommendation data, or some other type of data relevant to the operations of the CMO and/or marketing department.

[2361] In example embodiments, an executive digital twin, such as a CMO digital twin 60628 or other executive digital twins may depict a twin of a department, such as the marketing department or other department, which the user may use to identify, assign, instruct, oversee and review department personnel and third-party personnel that are associated with the activities of a particular department of an organization, including third-party partners and other outside associates involved

in the organization's related endeavors. Examples of such organization personnel include, but are not limited to, an organization's marketing staff, sales staff, finance staff, product design personnel, engineers, analysts, statisticians, data scientists, advertising staff, executive personnel, human resources staff, Board Members, advisors, or some other type of organization personnel. Examples of an organization's third-party personnel include, but are not limited to, advertising firm staff, ad exchange staff, outside creative or content developers, social media platform personnel, co-marketing partners, consultants, contractors, financial firm staff, auditors, or some other type of third-party personnel. In example embodiments, the departmental twin (in this example a marketing department twin) may include a definition of the various roles/employees working under the executive (e.g., CMO), the reporting structure, and associated permissions, for each individual in the business unit, and may be populated with the various names and/or other identifiers of the individuals filling the respective roles. In example embodiments, the department twin (e.g., marketing department twin) may include subsections that are specific to an activity or initiative, such as a marketing or advertising campaign. In this way, the executive (e.g., a CMO) may easily identify the personnel and third-party providers that are involved in the initiative and/or assign individuals and/or third parties to the initiative. A user may define one or more restrictions, permissions, and/or access rights of the individuals indicated in the business unit (e.g., using the configuration system 60002), as described herein, such that the restrictions, permissions, and/or access rights can be controlled by the CMO (or analogous user). In example embodiments, the permissions to define such restrictions and/or rights may be, for example, defined in the organizational digital twin that lists the user as having a role that permits implementing permissions, restrictions, and/or access rights to roles/individuals. In example embodiments, a personnel restriction or right associated with a role/individual may be specific to a project, such as a marketing or advertising campaign, and may define one or more types of data that a particular user or group of users is allowed, or not allowed, to access (either directly or in a digital twin). For example, a first marketing campaign twin may allow a marketing department employee to review the first marketing budget for a first marketing campaign and approve marketing expenditures for the first marketing campaign up to $10,000, but a second marketing campaign twin may disallow the same employee from any budgetary review or expenditures. Similar approaches can be used by projects of various types across an organization and its departments, such as product development projects, logistics projects, corporate development projects, service projects, and many others. In example embodiments, a breach, or attempted breach, of a restriction, permission or access right may invoke a notice, alert, warning or some other action to an individual notifying them of the breach or attempted breach. In examples, such a notice, alert, or warning may be sent to an individual that is identified based at least in part on the individual's position in the org chart relative to the person breaching or attempting to breach a restriction, permission or access right. In another example, such a notice, alert, or warning may be sent to an individual that is not identified in a departmental org chart and/or specific project or campaign, but rather may be sent to an individual that is identified based at least in part on a rule that is defined in the organizational twin of the entire enterprise. For example, a rule stored within an organizational digital twin of the entity

may specify that an alert must be sent to an Information Security Department staff member, or some other staff member, upon an attempted login to a forbidden file, or other, system. Other rules may be related to geographic, temporal, or other types of restrictions, as described herein. In example embodiments, an alert may be an email, phone call, text, or some other communication type.

[2362] In example embodiments, a CMO digital twin 60628 may be configured to oversee and manage personnel and human resources issues and activities related to the marketing department. For example, a marketing department twin may map each individual within the marketing department to her respective marketing department. Using the CMO digital twin 60628, the user may be able to select a department to see greater detail on the functioning of the department. Alternatively, this step may be automatically performed by the CMO digital twin 60628, requiring no action from the user (e.g., the CMO) (e.g., via an executive agent trained by the user). For example, the greater detail might include the number of vacancies currently associated with the department and the duration that each of the open positions has remained unfilled, estimated salary data associated with the open positions, and the like. The user may be able to also select to see more information on the budget associated with a given department, such as a department with a personnel vacancy, in order to see if there is currently available budget to cover a new hire for the department. Alternatively, this step may be automatically performed by the CMO digital twin 60628, requiring no action from the user. Continuing the example, if there is budget to cover a new hire, the CMO digital twin 60628 may provide a link or other opportunity for the user to initiate a communication with human resources or some other department personnel to begin the process of posting a job listing. Alternatively, this step may be automatically performed by the CMO digital twin 60628 (e.g., via an executive agent executing on behalf of the user), requiring no action from the user. This communication may be drawn from a repository of form emails, letters or other communications so that the user need not compose the communication, but rather only signal within the CMO digital twin 60628 that such communication should be sent. Similarly, based on the communication type (e.g., "initiate a new marketing job posting") the user may not need to select the receiving party, which may be stored in the EMP as the appropriate recipient based at least in part on a rule associated with the communication type. Continuing the example further, alternatively, if there is no budget available to cover a new hire, a second type of communication may be invoked by the CMO digital twin 60628, for example, an email, calendar invitation to reserve a meeting, or some other type of communication may be selected to be sent to the CFO, or other financial personnel, to request a meeting to discuss the marketing department's budget or initiate some other activity. Following this example, if and when the new hires are approved, the CMO digital twin may allow the user to delegate the hiring task to a subordinate or herself. In the event the user is assigned the hire the new employee, the CMO digital twin 60628 may provide materials regarding candidates (e.g., resume, referrals, interview notes from interviewers, or the like) and the user may select one or more candidates to further consider, interview, or hire.

[2363] In examples, a user may be able to select a sub-department within the marketing department to view the performance of the sub-department in greater detail. For example, the

greater detail might include the number of types of training sessions, tutorials, events, conferences, and the like that personnel in the selected marketing department have received. The user may be able to compare such training and event attendance levels with a specified target criterion that is stored in EMP, or that is associated with the EMP. This may result in the CMO digital twin 60628 reporting to the CMO a listing of personnel in her department whose training and/or event attendance fails to meet the target criterion. This listing may be prioritized by the CMO digital twin 60628 to highlight those staff members most in need of further training. The user may be able to also select to see more information on the budget associated with a given department, such as a department with staff who do not have adequate training according to the target criterion, in order to see if there is currently available budget to cover additional training for the department. If there is budget to cover additional training, the CMO digital twin 60628 may provide, for example, a link or other opportunity for the user to initiate a communication to a staff member in need of training to alert them that they must schedule training and/or attendance at an event within a timeframe. This communication may be drawn from a repository of form emails, letters or other communications so that the user need not compose the communication, but rather only signal within the CMO digital twin 60628 that such communication should be sent. Continuing the example further, a second type of communication may be invoked by the CMO digital twin 60628, for example, a request for information, training registration, or some other type of communication may be selected to be sent to a third-party training vendor that is used by the marketing department, a conference event registration, or other training or event entity, to request scheduling training and/or event registration, or some other activity. Alternatively, the steps, discussed above, for tracking and reporting on marketing personnel training and attendance may be automatically performed by the CMO digital twin 60628, requiring no action from the user. References to features and functions of the EMP and digital twins in this example of a marketing department and a CMO digital twin 60628 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2364] In example embodiments, a CMO digital twin 60628 may be configured to interface with the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the marketing department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein.

[2365] In example embodiments, a CMO digital twin 60628 may be configured to research, create, track and report on a marketing department budget including, but not limited to, an overall department budget, a budget for a single or group of marketing or advertising campaigns, a budget for a third-party vendor, or some other type of budget. The CMO digital twin 60628 may interact with and share such budget data and reporting with other executive twins, as described herein, including, but not limited to, a digital twin related to the finance department, accounts payable, executive staff such as the CEO and CFO, or others. The CMO digital twin 60628 may include

intelligence, based at least in part on the data analytics, machine learning and A.I. processes, as described herein, to read marketing budgets and related summaries and data in order to identify key departments, personnel, third-party or others that are, for example, listed in, or subject to, the budget line item and who therefore may have an interest in such material. Budget material pertaining to a given party may be abstracted and summarized for presentation independent from the entirety of the budget, and formatted and presented automatically, or at the direction of a user, to the party that is the subject of the budget item. In a simplified example, a CMO may create a new marketing campaign, "Airline – Airfare coupon texting campaign – January," which includes the following line items: Third-party advertising firm content creation $15,000; Social media platform placement $50,000; analytics department $25,000, and so forth. The entirety of the budget may be shared (at the election of the user or automatically) with parties that must approve the full budget, such as a CFO. As described herein this sharing may be accomplished by the CMO digital twin 60628 communicating directly with a CFO digital twin, so that the information is presented to the CFO without requiring the CFO to have knowledge of the budget or requesting the budget. Subparts of the budget, for example, the analytics department line item, may be automatically sent to the head of the analytics department by the CMO digital twin 60628 to inform that department of the total amount of authorized spending that is approved for that department for the specific marketing campaign.

[2366] In example embodiments, a CMO digital twin 60628 may be configured to track and report on inbound and outbound billing (i.e., accounts receivable and payable) related to the marketing department. The billing department, personnel, processes and systems, including a Billing digital twin may interact with the CMO digital twin 60628 to present, store, analyze, reconcile and/or report on billing activities related to parties with whom the marketing department is contracting, such as ad agencies, ad networks, ad exchanges, content creators, advertisers, social media platforms, television, radio, online entities, or others.

[2367] In example embodiments, a CMO digital twin 60628 may be configured to depict marketing campaign twins. In these embodiments, the CMO digital twin 60628 may depict various states and/or items relating to a marking campaign such as marketing content associated with a marketing campaign, market research performed with respect to a marketing campaign, tracking data of marketing content associated with marketing campaigns (e.g., geographic reach of marketing campaigns, demographic data associated with campaigns, etc.), analyses of marketing campaigns (e.g., outcomes related to marketing campaigns on various platforms), and the like. In example embodiments, a CMO digital twin may be configured to automatically report on marketing campaign-related activity via a user interface associated with the CMO digital twin 60628. Such activities may be determined using marketing department metadata that indicate state changes, such as an alteration to a website content, a change to a product photograph in an advertisement, a change in wording of a mailing, and the like. The CMO digital twin 60628 may also depict activity among a class of entities that are monitored or that are specified for monitoring in the CMO digital twin 60628, such as a new press release regarding a discounted advertising opportunity available from an ad exchange. In example embodiments, a CMO digital twin 60628

may be configured to provide research, tracking, monitoring, and analyses of media content performance across various marketing related platforms, and automatically report on such activity to a user interface associated with the CMO digital twin 60628. Such platforms may include, but are not limited to, customer relationship platforms (CRMs), organization website(s), social media, blogs, press releases, mailings, in-store or other promotions, or some other type of marketing platform-related material or activity.

[2368] In some of these embodiments, the CMO digital twin 60628 may be configured to simulate marketing campaigns, such that the simulations of the marketing campaign may vary parameters such as vehicles (e.g., social media, television, billboards, print, etc.), budget, targeting parameters (e.g., geographic, demographic, or the like), and/or other suitable marketing campaign parameters. In these embodiments, the digital twin simulation system 60320 may receive a request to perform the simulation CMO digital twin, where the request indicates campaign features and the parameters that are to be varied. In response, the digital twin simulation system 60320 may return the simulation results to the CMO digital twin 60628, which in turn outputs the results to the user via the client device display. In this way, the user is provided with various outcomes corresponding to different parameter configurations. In example embodiments, the user may select a parameter set based on the various outcomes. In example embodiments, an executive agent trained by the user may select the parameter sets based on the various outcomes.

[2369] In example embodiments, a CMO digital twin 60628 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to a marketing strategy, plan, campaign or initiative. For example, the CMO digital twin 60628 may be associated with a plurality of databases or other repositories of marketing presentation materials, summaries and reports and analytics, including such presentation materials, summaries and reports and analytics related to prior marketing campaigns, each of which may be further associated with financial and performance metrics pertaining to the campaign and which are also accessible to the CMO digital twin 60628. Such historical marketing campaign material may consist of advertising, marketing or other content that may be categorized based in part on the financial and performance metrics with which it is associated. For example, there may be a first category called "Market Tested Content," which consists of content that has been field deployed in a marketing campaign within a customer population, the actual performance of which is therefore fully known based on actual market testing. Because the marketing content from this category has been field tested, the content may be scored based at least in part on the financial, performance or other data with which it is associated. A second category may be "New Content – Simulation Tested," which consists of content that has not been deployed in the field, but which has been subject to analytic testing such as simulated customer segmentation analysis, simulated A/B testing, simulated attribution modeling, simulated market mix modeling, machine learning, A.I. techniques including, but not limited to, classification, probabilistic modeling, learning techniques, and the like. Because the marketing content from this category has been simulation tested, the content may be scored based at least in part on the simulated performance data or other data with which it is associated. Continuing the example, a third category of content may be "New Content – Panel Tested," which consists of

content that has not been deployed in the field, nor simulation tested, but which has been subject to testing among a human panel for their views, opinions and impressions. Because the marketing content from this category has been human panel tested, the content may be scored based at least in part on the performance data, as reported by the human panel, or other data with which it is associated. A final, fourth category of content may be "New – Untested," which is newly developed or other content that has not been tested in the field, in simulation, or by a human panel. The CMO digital twin 60628 may utilize the machine learning, A.I. and other analytic capabilities, as described herein, to analyze the content of the four categories of content and classify and score the content characteristics that are probabilistically associated with improved financial or other performance for stated types of marketing campaigns or marketing subject matter. Statistical weights may be applied to such characteristics, where the weight is indicative of a greater degree of financial or some performance metric of interest. Similarly, the characteristics of the market may be analyzed vis-a-vis the marketing content to determine the consumer characteristics that are probabilistically associated with improved financial or other performance for given marketing content. The CMO digital twin 60628 may provide a user interface within which access to this repository of stored data on content category, consumer and performance is available. When planning a marketing campaign, the CMO, or other marketing personnel, may use the CMO digital twin 60628 to select from this repository of content, that content which probabilistically will perform better with the intended consumer targets of the new campaign. For example, from historical marketing field tests from actual prior marketing campaigns, the data may show that marketing content having images of large dogs outperformed (based on, for example, ad conversion rates) content picturing small dogs, and this effect was positively correlated with age (i.e., older persons have an even greater preference for larger dogs). The performance data from the simulation-tested content may show a similar, but smaller effect based on the size of the dog images in the content, and the panel-tested data may show a similar effect for large dog imagery in content, but also have performance data indicating that the effect appears, based on the panel data, to be muted for persons 15 years or younger (i.e., young persons are more attracted to smaller dog breeds than older persons). For the CMO using the CMO digital twin 60628 this data, and the characteristics of the more successful content, may be used to select from the fourth category of content ("New – Untested") that content that is most appropriate for a new marketing campaign intended to sell a soft drink. In example embodiments, the artificial intelligence services system 60012 of the EMP 60000 may select the content and segment its presentation based at least in part on the prior performance data, so that the ads that are presented on platforms that tend to have persons over 15 will use content having a predominance of large breed dogs, and those platforms with younger audiences will offer a greater mix of dog breeds and possibly a preference for small breed dogs in marketing images. As the marketing campaign is deployed to the field, the CMO digital twin 60628 may monitor, track and report on the marketing campaign's performance so that the CMO can review and intervene as necessary. Once the new content has been field tested it may be stored and classified in the first category of content, "Market Tested Content," along with the related financial and performance metrics. In another example, similar stored content, content

categories, characteristics and financial and performance metrics may be used by the CMO digital twin 60628 to recommend, for example, search engine optimization (SEO), or other marketing strategies and techniques.

[2370] In example embodiments, a CMO digital twin 60628 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to market surveys, online surveys, customer panels, ratings, rankings, marketing trend data or other data related to marketing. A CMO digital twin 60628 may link to, interact with, and be associated with external data sources, and able to upload, download, aggregate external data sources, including with the EMP's internal data, and analyze such data, as described herein. Data analysis, machine learning, AI processing, and other analysis may be coordinated between the CMO digital twin 60628 and an analytics team based at least in part on using the artificial intelligence services system 60012. This cooperation and interaction may include assisting with seeding data elements and domains in the enterprise data store 60014 for use in modeling, machine learning, and AI processing to identify the optimal marketing content, sales channels, target consumers, price points, timing, or some other marketing-relating metric or aspect, as well as identification of the optimal data measurement parameters on which to base judgment of a marketing endeavor's success. Examples of data sources 60030 that may be connected to, associated with, and/or accessed from the CMO digital twin 60628 may include, but are not limited to, a sensor system 60032, a sales database 60034 that is updated with sales figures in real time, a CRM system 60038, a marketing campaign platform 60040, news websites, a financial database 60048 that tracks costs of the business, surveys 60050 (e.g., customer satisfaction surveys), an org chart 60052, a workflow management system 60054, customer databases 60062 structured to store customer data, and/or third-party datastores 60060 structured to store third-party data.

[2371] In example embodiments, a CMO digital twin 60628 may be configured to assist in the development of a new marketing campaign. For example, the CMO digital twin 60628 may identify an internal and external partner team for a marketing campaign. For example, individuals who are ideal candidates to assist with a marketing campaign may be identified based at least in part on experience and expertise data that is stored within or in association with the CMO digital twin 60628. In another example, the CMO digital twin 60628 may identify marketing campaign goals and record, monitor and track the campaign's performance relative to those goals and present, in real-time, the tracking of the campaign to the CMO within a user interface that is associated with the CMO digital twin 60628. Examples of marketing targets include, but are not limited to, unit distribution, customer acquisition customer retention, customer churn, customer loyalty (e.g., repeat purchases), customer acquisition costs, duration of an average sales cycle, ad conversion rate, sales growth, geographic expansion of sales, demographic expansion of sales, market penetration, percentage of market control, marketing campaign ROI, regional comparison of performance, channel analysis, sales partner analysis, marketing partner analysis, or some other marketing target.

[2372] In example embodiments, a CMO digital twin 60628 may be configured to monitor customer feedback loops, customer opinions, customer satisfaction, complaints, product returns

and the like based at least in part on use of the monitoring agent of the client application 60104, as described herein, that is associated with the CMO digital twin 60628. Such feedback data may include, but is not limited to, data that derives from call center activity, chatbot activity, email (e.g., complaints), product returns, Better Business Bureau submissions, or some other type of customer feedback or manifestation of customer opinion. The client application 60104 may include a monitoring agent that monitors the manner by which customers or others respond to a marketing campaign. The monitoring agent may report the customer's response to such campaigns to the EMP 60000 for presentation in a user interface that is associated with the CMO digital twin 60628. In response, the EMP 60000 may train an executive agent (which may include one or more machine-learned models) to handle and process such notifications when they next arrive, and escalate and/or alert the CMO when such notifications are of an urgent nature, for example, an announcement of a class action lawsuit related to a product that is the subject of a marketing campaign. In example embodiments, the CMO digital twin 60628 may generate performance alerts based on performance trends. This may allow a CMO to optimize marketing campaigns in real-time without having to manually request such real-time performance data; the CMO digital twin 60628 may automatically present such information and related/necessary alerts as configured by the organization, CMO, or some other interested party.

[2373] In example embodiments, a CMO digital twin 60628 may be configured to report on the performance of the marketing department, personnel of the marketing department, marketing campaigns, marketing content, marketing platforms, marketing partners, or some other aspect of management within a CMO's purview. Reporting may be to the CMO, the marketing department, to other executives of an organization (e.g., the CEO), or to outside third parties (e.g., marketing partners, press releases, and the like). As described herein, reporting may include sales summaries, customer data, marketing campaign performance metrics, cost-per-sale data, cost-per-conversion data, customer analysis, such as predicted customer lifetime value for newly acquired customers, or some other type of reporting data. Reporting and the content of reporting may be shared by the CMO digital twin 60628 with other executive digital twins, for example, data related to new customers having a particularly high predicted customer lifetime value may be shared with a sales staff for the purpose of exploring cross-selling opportunities. The reporting functionality of the CMO digital twin 60628 may also be used for populating required data for formal reporting requirements such as shareholder statements, annual reports, SEC filings, and the like. Templets of common reporting formats may be stored and associated with the CMO digital twin 60628 to automate the presentation of data and analytics according to pre-defined formats, styles and system requirements.

[2374] In example embodiments, a CMO digital twin 60628 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to competitors of a CMO's organization, or named entities of interest. In example embodiments, such data may be collected by the EMP 60000 via data aggregation, spidering, web-scraping, or other techniques to search and collect competitor information from sources including, but not limited to, press releases,

SEC or other financial reports, mergers and acquisitions activity, or some other publicly available data.

[2375] In example embodiments, a CMO digital twin 60628 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to regulatory activity, such as government regulations, industry best practices or some other requirement or standard. For example, the marketing industry is subject to data privacy and security laws in many jurisdictions, and it is an area of law and regulation that is experiencing rapid change. In example embodiments, the CMO digital twin 60628 may be in communication with another enterprise digital twin, such as a General Counsel digital twin 60634, through which the legal team can keep the CMO apprised of new regulation or regulation changes as they occur. Similarly, as a CMO develops new market campaigns and selects the jurisdictions (e.g., the United States vs Europe) and populations that will be a part of the campaigns (e.g., minors vs. adults), the CMO digital twin 60628 may automatically send a synopsis of the aspects of the campaigns that are relevant for privacy law review so that the campaign may be vetted for legal and regulatory compliance prior to launch. In examples, such a marketing campaign synopsis might include a summary of the jurisdictions of the campaign, intended audience, means of obtaining consent, the type of consent to be obtained (e.g., opt-in, opt-out, passive), and so forth. Once approved and launched, as customer consents and other data privacy-related information is received by an organization, the CMO digital twin 60628 may facilitate the CMO tracking metrics, for example the percentage of customers choosing to opt-in to receive future marketing material (e.g., email solicitations). As the organization receives privacy related material it may store such information for future retrieval, summary, deletion or other activity, for example, in response to a data subject request from an EU citizen who has requested their data be deleted (i.e., exercising their "right to be forgotten"). In example embodiments, the CMO digital twin 60628 may monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to what customer data is collected, the party responsible for its collection and storage, the location and duration of storage, and so forth. This data may be called forth by the CMO digital twin 60628, for example, in the event of a data breach. The CMO digital twin 60628 may be able to summarize, for example, a list of persons affected by the breach and the type of data that was breached and share this information with a Chief Privacy Officer (CPO), including sharing with the CPO digital twin.

[2376] In example embodiments, the client application 60104 that executes the CMO digital twin 60628 may be configured with an executive agent that reports a CMO's behaviors and preferences (or other marketing personnel's behaviors and preferences) to the expert agent system 60010, as described herein, and the expert agent system 60010 may train the executive agent on how the CMO or other marketing personnel respond to certain situations and adjust its operation based at least in part on the data collection, analysis, machine learning and A.I. techniques, as described herein.

[2377] In example embodiments, a Chief Technical Officer (CTO) digital twin 60630 may be a digital twin configured for a CTO or other technology executive of an enterprise tasked with overseeing and managing the R&D, technology development, technical implementations of the

enterprise, and/or engineering activities of the enterprise. In example embodiments, a CTO digital twin 60630 provides real-time views of enterprise technology assets, including technology capabilities and versions. For example, in a manufacturing enterprise, a CTO digital twin 60630 may depict where environment-compatible updates, upgrades, or substitutions may be available. A CTO digital twin 60630 may provide data, analytics, summary, and/or technical reporting including, but not limited to, real-time, historical, aggregated, comparison, and/or forecasted technical information (e.g., real-time, historical, simulated, and/or forecasted technical performance data related to company products, benchmarking results, and the like). By using a CTO digital twin 60630, a CTO may be better able to stay abreast of technical developments and software engineering impacts by engaging in continuous virtualized learning using the CTO digital twin 60630. In example embodiments, a CTO digital twin 60630 may assist in virtual collaboration (a CTO-essential skill), as a CTO will need to partner with in-house engineers and external vendors in a virtual environment to imagine and ideate to achieve something, often something that has not been done before. In example embodiments, the CTO digital twin may work in connection with the EMP 60000 to provide simulations, predictions, statistical summaries, decision support based on analytics, machine learning, and/or other AI and learning-type processing of inputs (e.g., technical performance data, sensor data and the like).

[2378] In example embodiments, a CTO digital twin 60630 may provide features and functionality including, but not limited to, management of technical personnel, partners and outside consultants and contractors (e.g., developers, beta testers, and the like), oversight of budgets, procurement, expenditures, policy compliance (e.g., policies related to code usage, storage, documentation, and the like), and other technology, development, and/or engineering-related resources, and/or reporting.

[2379] In example embodiments, the types of data that may populate a CTO digital twin may include, but are not limited to, technology performance and specification data, interoperability and compatibility data, cybersecurity data, competitor data, failure mode effects analysis (FMEA) data, technology/engineering roadmap data, information technology systems data (including with respect to any of the hardware, software, networking, and other types mentioned or described herein), operations technology and systems data, uptime/downtime/operational performance data, asset aging/vintage/timing data, technical performance metrics by business unit, by product, by geography, by factory, by store location(s), resource utilization, competitive product and pricing data, forecast data, demand planning data, analytic results of AI and/or machine learning modeling (e.g., technical forecasting), prediction data, metrics relating to patent disclosures, patent filings, and/or patent grants, recommendation data, and/or other types of data relevant to the operations of the CTO and/or technology, development, and/or engineering department.

[2380] In example embodiments, a CTO digital twin 60630 may depict a twin of a set of technology, development, and/or engineering departments, which the user may use to identify, assign, instruct, oversee and review technology, development, and/or engineering department personnel and third-party personnel that are associated with the technology, development, and/or engineering activities of an organization, including third-party partners and other outside

contractors, such as third-party developers and/or testers that are involved in the organization's technology, development, and/or engineering activities. Examples of such organization personnel include, but are not limited to, technology, development, and/or engineering department staff, sales staff and analysts, statisticians, data scientists, or some other type of organization personnel relevant to the functioning of a technology, development, and/or engineering department. Examples of a technology, development, and/or engineering department's third-party personnel include, but are not limited to, management consultants, developers, software engineers, testers, and/or engineering partners, consultants, contractors, technical firm staff, auditors, or some other type of third-party personnel.

[2381] In example embodiments, the CTO digital twin 60630 may include a definition of the various roles/employees working under the CTO, the reporting structure, and associated permissions, for each individual in the business unit, and may be populated with the various names and/or other identifiers of the individuals filling the respective roles.

[2382] In example embodiments, a client application 60104 executing a CTO digital twin 60630 may interface with the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the technology, development, and/or engineering department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein. Collaboration and communication tools and associated rules may be configured to use company-, industry- and domain-specific taxonomies and lexicons when representing entities, states and flows within the CTO digital twin 60630.

[2383] In example embodiments, a CTO digital twin 60630 may be configured to allow a user to research, create, track and report on a technology, development, and/or technology or engineering department initiative including, but not limited to, a new product development, update, enhancement, replacement, upgrade, or the like. In example embodiments, the CTO digital twin 60630 may be associated and/or in communication with databases, including databases storing analytic and/or product data and product performance data, and present information to an interface associated with the CTO digital twin 60630, as described herein. As product development advances, real time operations and other technical information may be used to continuously update the product development summary that is available for the CTO or other technical personnel to review. The CTO digital twin 60630 may also be associated and/or in communication with databases, including databases storing analytic and/or competitive product data and product performance data, and present this information to an interface associated with the CTO digital twin 60630, as described herein. As the CTO's company's products change, and competitor products change, their current state and specifications may be presented by the CTO digital twin 60630 for the CTO or other technical personnel to review direct product comparisons. Such comparisons may be used, in part, to produce analytics, scores, reports and the like indicating the relative advantages and/or disadvantages that a company's product(s) has relative to competitor product(s). In

examples, a report may be automatically provided to the marketing department to emphasize the relative advantages that a company product has over a competitor product (e.g., speed of processing) that should be used in a new marketing campaign. Sharing with the marketing department may be accomplished, in part, by the CTO digital twin 60630 communicating with the CMO digital twin 60628 to present reports or other information to the CMO or marketing staff.

[2384] In example embodiments, the CTO digital twin 60630 may be configured to present simulations of technology development and/or engineering activities. For example, in example embodiments, the digital twin system 60004 may simulate product usage under a plurality of constraints that might impact product performance, such as an operating environment, processing speed, storage or other platform characteristics. In example embodiments, real time operations data, such as operations data available through the EMP 100, may be incorporated into simulated data for the purposes of running operational simulations. This may allow a CTO to a gain a deeper understanding of the operation of the company's products in the real world and within an altered, simulated real world environment. It may also allow operational digital twin-based product architectures to be built that link actual product production with business priorities to enable simulated decision making in a virtual environment and assist in the evaluation of vendor supplied solutions by enabling the review of such digital twins in the context of their supplied solutions and the relationship to the business. In example embodiments, simulations may also include simulations related to varying technical and/or product specification parameters, product design and monitoring, internal controls design, testing, certification, and deliver technical and non-technical data in reports, presentations, and dashboards for technical decision making. In these embodiments, the digital twin simulation system 60320 may receive a request to perform the simulation requested by the CTO digital twin 60630, where the request indicates features and the parameters, including technical parameters, that are to be varied. In response, the digital twin simulation system 60320 may return the simulation results to the CTO digital twin 60630, which in turn outputs the results to the user via the client device display. In this way, the user is provided with various outcomes corresponding to different technical and/or product parameter configurations. In example embodiments, the user may select a parameter set based on the various outcomes. In example embodiments, an executive agent trained by the user may select a technical parameter set based on the various outcomes. The simulations, analytics and/or modeling performed by the CTO digital twin 60630 may be used to reduce testing time, design time, or some other type of technical cost. The simulations, analytics and/or modeling performed by the CTO digital twin 60630 may be used to create and structure product development and testing plans. The simulations, analytics and/or modeling performed by the CTO digital twin 60630 may be used to evaluate product go-to-market timing and preparedness. The CTO equipped with digital twin 60630 will be better able to adapt quickly to identify product and/or technical parameters in need of further development and predict products' operational performance. This may reduce errors, speed testing and reduce the need for patches, bug fixes, updates and the like and flatten agile process management.

[2385] In example embodiments, a CTO digital twin 60630 may provide an interface that allows a user to research, create, track and report on a technology, development, and/or engineering department initiative including, but not limited to, an overall department budget, a budget for a single or group of technology, development, and/or engineering initiatives, a third-party vendor activity, or some other type of expense or budget. The CTO digital twin 60630 may interact with and share such expense or budget data and reporting with other executive twins, including, but not limited to, a digital twin related to accounts payable, executive staff such as the CEO, and/or others.

[2386] In example embodiments, the CTO digital twin 60630 may leverage the artificial intelligence services system 60012 (e.g., data analytics, machine learning and A.I. processes) to read technical reports, projections, simulations, and related summaries and data in order to identify key departments, personnel, third-party or others that are, for example, listed in, or subject to, a technical item or detail provided.

[2387] In example embodiments, a CTO digital twin 60630 may be configured to provide a CTO, or other technology, development, and/or engineering department personnel, with information that is unique to the CTO digital twin 60630 and thus can provide insights and perspectives on technical performance that are unique to the CTO digital twin 60630, based at least in part on the CTO digital twin 60630 make making use of real time production, development and operational data based on both real world and simulated activity.

[2388] In example embodiments, the CTO digital twin 60630 may be configured to obtain and depict oversight activity that includes, but is not limited to, internal controls design, testing, and reporting while directing listed actions the appropriate personnel.

[2389] In example embodiments, a CTO digital twin 60630 may be configured to depict, aggregate, merge, analyze, prepare, report and distribute material relating to a technical strategy, plan, activity or initiative. For example, the CTO digital twin 60630 may be associated with a plurality of databases or other repositories of technical materials, summaries and reports and analytics, including such materials, summaries and reports and analytics related to prior technical activity and results (e.g., bug testing), each of which may be further associated with third-party technical or economic data, including competitor product data and/or technical benchmarks.

[2390] In example embodiments, a CTO digital twin 60630 may be configured to depict, aggregate, merge, analyze, prepare, report and distribute material relating to technical reporting, ratings, rankings, technical trend data, or other data related to company technology, development, and/or engineering. A CTO digital twin 60630 may link to, interact with, and be associated with external data sources, and able to upload, download, aggregate external data sources, including with the EMP's internal data, and analyze such data, as described herein. Data analysis, machine learning, AI processing, and other analysis may be coordinated between the CTO digital twin 60630 and an analytics team based at least in part on using the intelligence services system 60012. This cooperation and interaction may include assisting with seeding technology, development, and/or engineering-related data elements and domains in the enterprise data store 60014 for use in modeling, machine learning, and AI processing to identify the optimal technical strategy, or some other technology, development, and/or engineering-relating metric or aspect, as well as

identification of the optimal data measurement parameters on which to base judgment of a technology initiative, development initiative, and/or engineering endeavor's success. Examples of data sources 60030 that may be connected to, associated with, and/or accessed from the CTO digital twin 60630 may include, but are not limited to, a sensor system 60032, a sales database 60034 that is updated with sales figures in real time, a technology, development, and/or engineering platform 60040, news websites, a technical database that tracks costs of the business, an org chart 60052, a workflow management system 60054, customer databases 60062 that store customer data, and/or third-party datastores 60060 that store third-party data.

[2391] In example embodiments, a CTO digital twin 60630 may aggregate data sources and types, creating new data types, summaries and reports that are not available elsewhere. This may reduce reliance upon the need of multiple third-party providers and current solutions. This may, among other benefits and improvements, reduce expenses associated with acquiring data needed for sound technical decision making.

[2392] In example embodiments, a CTO digital twin 60630 may be configured to monitor technical performance, including real time monitoring, based at least in part on use of the monitoring agent of the client application 60104, as described herein, that is associated with the CTO digital twin 60630. The monitoring agent may report on such activities to the EMP 60000 for presentation in a user interface that is associated with the CTO digital twin 60630. In response, the EMP 60000 may train an executive agent (which may include one or more machine-learned models) to handle and process such notifications when they next arrive, and escalate and/or alert the CTO when such notifications are of an urgent nature, for example, an identification of a new technical bug or a security patch that is urgently needed. In example embodiments, the CTO digital twin 60630 may generate technical performance alerts based on performance trends. This may allow a CTO to optimize initiatives in real-time without having to manually request such real-time technical performance data; the CTO digital twin 60630 may automatically present such information and related/necessary alerts as configured by the organization, CTO, or some other interested party.

[2393] In example embodiments, a CTO digital twin 60630 may be configured to report on the performance of the technology, development, and/or engineering department, personnel of the technology, development, and/or engineering department, technology, development, and/or engineering activities, technology, development, and/or engineering content, technology, development, and/or engineering platforms, technology, development, and/or engineering partners, or some other aspect of management within a CTO's responsibilities. Reporting may be to the CEO, the technology, development, and/or engineering department, to other executives of an organization (e.g., the CIO), or to outside third parties.

[2394] In example embodiments, a CTO digital twin 60630 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to industry best practices, benchmarks, or some other requirement or standard. For example, the CTO digital twin 60630 may be in communication with another enterprise digital twin, such as a CIO digital twin 60632, through which the technical team can keep the CIO apprised of changes as they occur.

[2395] In example embodiments, the client application 60104 that executes the CTO digital twin 60630 may be configured with an executive agent that reports a CTO's behaviors and preferences (or other technology, development, and/or engineering personnel's behaviors and preferences) to the executive agent system 60010, as described herein, and the executive agent system 60010 may train the executive agent on how the CTO or other technology, development, and/or engineering personnel respond to certain situations and adjust its operation based at least in part on the data collection, analysis, machine learning and A.I. techniques, as described herein.

[2396] References to features and functions of the EMP and digital twins in this example of a CTO digital twin 60630 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2397] In example embodiments, a Chief Information Officer (CIO) digital twin 60632 may be a digital twin configured for the CIO of an enterprise, or analogous executive tasked with overseeing the intelligence, information, data, knowledge, and/or IT operations of the enterprise. In example embodiments, a CIO digital twin 60632 depicts a real time representation of an organization's information assets and workflows including data relating to data security, network security and enterprise knowledge. The real time representation may be based at least in part on real-time operations data that tracks the performance of an organization's information infrastructure, including internal information assets, customer-facing technologies, and information assets provided and/or serviced by third parties, such as cloud computing service providers. For example, a CIO digital twin 60632 may receive real time information regarding the performance of a network, such as an intranet used by an organization, APIs that are accessed by the enterprise, APIs that are exposed by the enterprise, software that is running on the enterprise's systems, or the like. The information may be aggregated and presented to a CIO in order to provide an overview of the general performance of the computing infrastructure of the enterprise. For example, the CIO digital twin may indicate whether there are any network outages occurring, whether there are any security risks detected in the enterprise's network, whether any software systems are operating improperly, and many other scenarios. In example embodiments, the CIO digital twin 60632 may present a user interface that allows a user (e.g., the CIO) to select particular network assets to review in greater detail, such as an asset the real time operations data indicates is experiencing an operational failure or other issues. Such real time operations data related to IT and other information asset performance may allow the CIO to better track the performance and needs of an organization's information and IT infrastructure and better enable troubleshooting issues, simulate solutions, select appropriate information and IT management actions, and maintain the organization's information and IT infrastructure.

[2398] In example embodiments, a CIO digital twin 60632 may provide data, analytics, summary, and/or information and IT reporting including, but not limited to, real-time, historical, aggregated, comparison, and/or forecasted information (e.g., real-time, historical, simulated, and/or forecasted performance data related to company information and IT assets, third-party assets, and the like). A CIO empowered by a CIO digital twin 60632 may be better able to maintain and evolve information and IT assets through continuous monitoring using the CIO digital twin 60632. A CIO digital twin

60632 may assist in virtual monitoring and testing in a virtual environment to test implementations, changes, reconfigurations, the introduction and/or removal of components and other assets, and the like. In example embodiments, the CIO digital twin may work in connection with the EMP 60000 to provide simulations, predictions, statistical summaries, decision support based on analytics, machine learning, and/or other AI and learning-type processing of inputs (e.g., performance data, sensor data, and the like).

[2399] In example embodiments, the types of data that may populate a CIO digital twin 60632 may include, but are not limited to, information and IT asset performance and specification data, interoperability and compatibility data, cybersecurity data, uptime/downtime/operational performance data, asset aging/vintage/timing data, resource utilization, results of AI and/or machine learning modeling (e.g., IT performance simulations), or some other type of data relevant to the operations of the CIO.

[2400] In example embodiments, a CIO digital twin 60632 may be configured to interface with the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the technology, development, and/or engineering department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein. Collaboration and communication tools and associated rules may be configured to use company-, industry- and domain-specific taxonomies and lexicons when representing entities, states and flows within the CIO digital twin 60632.

[2401] In example embodiments, the CIO digital twin 60632 may be configured to provide simulations of an organization's information and IT activities including, but not limited to network utilization, disaster planning, IT asset selection, maintenance protocols, downtime planning, and the like that is simulated under a plurality of hypothetical IT environments and scenarios that might impact performance, such as a security breach, IT asset failure, information failure, network congestion, or other activity or event. Real time operations data, such as that available through the EMP, as described herein, may be incorporated into simulated information or IT Infrastructure scenarios for the purposes of running operational simulations. The simulations, analytics and/or modeling performed by the EMP 100 with respect to a CIO digital twin 60632 may be used to reduce testing time, design time, or some other type of IT cost. The simulations, analytics and/or modeling performed by the CIO digital twin 60632 may be used to create and structure IT assets, networks, and guide development and testing plans. The simulations, analytics and/or modeling performed by the CIO digital twin 60632 may be used to evaluate network security, performance, and other features. The CIO equipped with digital twin 60632 may quickly identify optimal asset configurations to maximize operational performance.

[2402] In example embodiments, a CIO digital twin 60632 may be configured to provide a user (e.g., the CIO) with information that is unique to the CIO digital twin 60632 and thus can provide insights and perspectives on information and IT asset performance that are unique to the CIO digital twin 60632, based at least in part on the CIO digital twin 60632 make making use of real

time production, development and operational data based on both real world and simulated activity. In example embodiments, the CIO digital twin 60632 may be configured to manage operational planning, based at least in part by leveraging predictive analytics for development planning. In example embodiments, a CIO digital twin 60632 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to an information and/or an IT strategy, scenario, event, plan, activity or initiative. For example, the CIO digital twin 60632 may be associated with a plurality of databases or other repositories of information, materials, summaries and reports and analytics, including such materials, summaries and reports and analytics related to prior events, activity and results (e.g., a system outage).

[2403] In example embodiments, a CIO digital twin 60632 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to information and/or IT reporting, ratings, rankings, information, knowledge and IT trend data, or other data related to company information and/or IT assets and infrastructure. A CIO digital twin 60632 may link to, interact with, and be associated with external data sources, such that CIO digital twin 60632 may upload, download, aggregate external data sources, and/or analyze such enterprise data.

[2404] In example embodiments, a CIO digital twin 60632 may be configured to monitor IT performance, including in real time, based at least in part on use of the monitoring agent of the client application 60104, as described herein, that is associated with the CIO digital twin 60632. The monitoring agent may report on such activities to the EMP 60000 for presentation in a user interface that is associated with the CIO digital twin 60632. In response, the EMP 60000 may train an executive agent (which may include one or more machine-learned models) to handle and process such notifications when they next arrive and escalate and/or alert the CIO when such notifications are urgent.

[2405] In example embodiments, a CIO digital twin 60632 may be configured to report on the performance of an organization's IT assets, network, or some other aspect of management within a CIO's responsibilities. In example embodiments, the client application 60104 that executes the CIO digital twin 60632 may be configured with an executive agent that reports a CIO's behaviors and preferences to the executive agent system 60010, and the executive agent system 60010 may train the executive agent on how the CIO or other personnel respond to certain IT situations and adjust its operation based at least in part on the data collection, analysis, machine learning and A.I. techniques described throughout the disclosure.

[2406] References to features and functions of the EMP and digital twins in this example of a marketing department and a CIO digital twin 60632 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2407] In example embodiments, a general counsel (GC) digital twin 60634 may be an executive digital twin configured for the general counsel (GC) of an enterprise, or an analogous executive tasked with overseeing the legal department and/or outside counsel of the enterprise. A GC digital twin 60634 may provide functionality including, but not limited to, management of legal personnel, partners and outside counsel, oversight of legal budgets and resources, compliance, management

of contracting and litigation, management of internal policies, intellectual property, employment law, tax law, privacy law, reporting, and regulatory analysis.

[2408] In example embodiments, the types of data that may populate and/or be utilized by a GC digital twin 60634 may include, but are not limited to, budgetary data (e.g., external legal spend, internal legal spend, ancillary legal costs, and the like), regulatory data (e.g., regulatory requirements, regulatory actions taken, and the like); contract and licensing data (e.g., in progress negotiations, current contract obligations, past contract obligations, and the like); compliance data (e.g., compliance requirements, compliance actions taken, and the like), litigation data (e.g., potential litigations sources, pending litigations, past litigations, settlement agreements, and the like), employment data (e.g., employment contracts, employee complaints, employee stock options, and the like), intellectual property data (e.g., filed patent applications, patent dockets, issued patents, trademark applications, trademark docket data, registered trademarks, and the like), tax data, privacy data, regulatory data, analytic results of AI and/or machine learning modeling; prediction data; recommendation data, or some other type of data relevant to the operations of the GC and/or legal department.

[2409] In example embodiments, a GC digital twin 60634 may be configured based at least in part on using the collaboration suite 60008 to specify and provide a set of collaboration tools that may be leveraged by the legal department and associated parties. The collaboration tools may include video conferencing tools, "in-twin" collaboration tools, whiteboard tools, presentation tools, word processing tools, spreadsheet tools, and the like, as described herein. Collaboration and communication rules may be configured based at least in part on using the AI reporting tool, as described herein. Collaboration and communication tools and associated rules may be configured to use company-, industry- and domain-specific taxonomies and lexicons when representing entities, states and flows within the GC digital twin 60634, such as ones related to particular bodies of law, regulation, jurisdiction, or practice area, such as ones related to corporate law, commercial law, bankruptcy law, the law of secured transactions, banking law, customs law, export control regulations, maritime law, trade law, international treaties, securities law, contracts law, environmental law, international law, privacy law, data privacy law, patent law, civil and criminal procedure, trademark law, copyright law, trade secret law, unfair competition law, the law of torts, property law, advertising law, and many others.

[2410] In example embodiments, a GC digital twin 60634 may be configured to research, create, track and issue reports on a legal department budget including, but not limited to, an overall department budget, a budget for a specific project, such as "U.S. patent filings," or group of projects, a budget for a specific litigation, a budget for a third-party vendor, such as outside counsel, or some other type of legal budget. A GC digital twin 60634 may be configured to create, track, provide research, and report on financial data related to material under review or supervisions of the legal department including, but not limited to, licensing revenues, licensing expenditures, or some other type of financial data related to legal department review and responsibilities. In example embodiments, the GC digital twin 60634 may interact with and share such licensing revenue and/or budget data and reporting with other executive twins, as described herein,

including, but not limited to, a CFO digital twin 60622, CEO digital twin, COO digital twin, CTO digital twin, and the like. In example embodiments, the GC digital twin 60634 may include intelligence, based at least in part on the data analytics, machine learning and A.I. processes, as described herein, to read legal contracts, licenses, budgets and related summaries and data in order to identify key departments, personnel, third-party or others that are, for example, listed in, or subject to, or impacted by a license and/or budget line item and who therefore may have an interest in such material. License and/or budget material pertaining to a given party may be abstracted and summarized for presentation independent from the entirety of the budget, and formatted and presented automatically, or at the direction of a user, to the party that is the subject of the budget item. In a simplified example, a GC may have license(s) under her department's review which have line items, schedules, appendices and the like detailing licensing revenues that will be owed to the organization over a prescribed timeframe. The GC may use the GC digital twin 60634 to consolidate, summarize and/or share such financial data derived, or to be derived, from licensing revenues with another executive in an organization, such as the CFO (e.g., via a CFO digital twin) and/or CEO (e.g., via a CEO digital twin). The data shared may indicate the licensing revenues to be obtained in a given financial quarter to assist the CFO and others in maintaining an accurate and current summary of projected quarterly revenues.

[2411] In example embodiments, a GC digital twin 60634 may be configured to track and report on inbound (e.g., settlement or litigation revenue) and outbound billing (e.g., outside counsel costs) related to the legal department. The billing department, personnel, processes and systems may interact with the GC digital twin 60634 to present, store, analyze, reconcile and/or report on billing activities related to parties with whom the legal department is contracting, such as outside counsel, consultants, research services, online entities, or others. In example embodiments, a GC digital twin 60634 may be configured to research, track, monitor, store, analyze, create and distribute legal content, and automatically report on such activity to a user interface associated with the GC digital twin 60634. Such activities might include storing data so that the GC digital twin 60634 may detect a state change, for example, a new court filing in a litigation, a communication received from outside counsel, a new license draft from opposing counsel, a draft patent application, a notice from the United States Patent and Trademark Office, or some other type of new or updated material. The GC digital twin 60634 may also detect activity among a class of entities that are monitored or that are specified for monitoring in the GC digital twin 60634, such as particular courts, regulatory or legislative bodies or some other type of entity. In example embodiments, a GC digital twin 60634 may be configured to research, track, monitor, store, and analyze content of various legal related platforms, and automatically report on such activity to a user interface associated with the GC digital twin 60634. Such platforms may include, but are not limited to, bar or other legal associations, courts, legal search platforms, social media, legal blogs, press releases, or some other type of legal platform-related material or activity.

[2412] In example embodiments, a GC digital twin 60634 may be configured to store, aggregate, merge, analyze, prepare, report and distribute material relating to a legal strategy, legal documents, litigation, legal recommendations or some other legal activity. For example, the GC digital twin

60634 may be associated with a plurality of databases or other repositories of legal materials, contracts, licenses, intellectual property (e.g., patent filings), summaries and reports and analytics. A GC digital twin 60634 may link to, interact with, and be associated with external data sources, and able to upload, download, aggregate external data sources, including with the EMP's internal data, and analyze such data, as described herein. Data analysis, machine learning, AI processing, and other analysis may be coordinated between the GC digital twin 60634 and an analytics team based at least in part on using the intelligence services system 60012. This cooperation and interaction may include assisting with seeding data elements and domains in the enterprise data store 60014 for use in modeling, machine learning, and AI processing to identify the optimal and/or relevant legal content, legal documents, parties associated with a legal activity (e.g., a litigation), as well as identification of the optimal data measurement parameters on which to base judgment of a legal endeavor's success (e.g., licensing revenue, staying within a stated budget for the use of outside counsel, and the like). Examples of data sources 60030 that may be connected to, associated with, and/or accessed from the GC digital twin 60634 may include, but are not limited to, a legal research platform, legal websites, news websites, a financial database 60048, contracts database, one or more org charts 60052, a workflow management system 60054, and/or third-party datastores 60060 that store third-party data.

[2413] In example embodiments, a GC digital twin 60634 may be configured to assist in the development of a new legal endeavor, such as pursuit of a new contract, review of a new law or regulation impacting a business, litigation or arbitration, or some other legal activity. For example, the GC digital twin 60634 may identify an internal and external partner (e.g., outside counsel) team for a legal action. For example, individuals who are ideal candidates to assist with a legal action may be identified based at least in part on experience and expertise data that is stored within or in association with the GC digital twin 60634. For example, the GC may be initiating negotiations of a joint development agreement between entities that are located in the United States and Taiwan and may need to obtain outside Taiwanese counsel. Using the GC digital twin 60634, the GC may be presented with details of prior outside counsel used in Taiwan for similar projects. In another example, if the GC digital twin 60634 does not locate details of prior outside counsel used in Taiwan for similar projects, the GC digital twin 60634 may scan, research, collect and summarize information from public or other sources on highly rated, recommended or other Taiwanese outside counsel that may be appropriate, based on skills, experience and the like, to work on the joint development agreement project.

[2414] In example embodiments, the GC digital twin 60634 may identify legal project goals and record, monitor and track the project's performance relative to those goals and present, in real-time, the tracking of the project to the GC within a user interface that is associated with the GC digital twin 60634. For example, the GC digital twin 60634 may include a clickable dashboard that, when clicked, illustrates the status of a set of legal projects. In example embodiments, the dashboard may include timelines for each project and a relative status of each project with respect to its timeline.

[2415] In example embodiments, a GC digital twin 60634 may be configured to report on the performance of the legal department, personnel of the legal department, legal actions, legal content, legal platforms, legal partners, or some other aspect of a GC's management. Reporting may be to the GC, the legal department, to other executives of an organization (e.g., the CEO), or to outside third parties (e.g., outside counsel, legal notices, press releases, and the like). Reporting and the content of reporting may be shared by the GC digital twin 60634 with other executive digital twins, for example, data related to regulation compliance, ongoing litigation, or some other legal activity. The reporting functionality of the GC digital twin 60634 may also be used for populating required data for formal reporting requirements such as shareholder statements, annual reports, SEC filings, and the like. Templates of common reporting formats may be stored and associated with the GC digital twin 60634 to automate the presentation of data and analytics according to pre-defined formats, styles and system requirements. In example embodiments, the GC digital twin may be configured to leverage an executive agent 60704 trained on behalf of the GC to create and disseminate the reports.

[2416] In example embodiments, a GC digital twin 60634 may be configured to monitor, store, aggregate, merge, analyze, prepare, report and distribute material relating to regulatory activity, such as government regulations, regulatory compliance, legislation, court opinions, industry best practices or some other requirement or standard. For example, the GC digital twin 60634 may keep the GC apprised of new regulation or regulation changes as they occur. The GC may set parameters of the GC digital twin 60634 regarding the legal domains, subject matter areas, jurisdictions, or some other parameters, that are of interest to the GC that the GC digital twin 60634 should monitor.

[2417] In example embodiments, a GC digital twin 60634 may leverage an executive agent 60704 that is trained on user's (e.g., GC) behaviors and preferences (or other legal personnel's behaviors and preferences). In example embodiments, the client application 60104 hosting the GC digital twin 60634 may track the user's actions relating to various events, notifications, alerts, or the like and may report the tracked events using the expert agent system 60010, as described herein. In response, the expert agent system 60010 may learn how the GC or other legal personnel respond to certain situations and may train an execute agent 60704 on behalf of the user (e.g., GC), such that the executive agent 60704 may respond to similar situations once deployed.

[2418] References to features and functions of the EMP and digital twins in this example of a legal department and a GC digital twin 60634 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2419] In example embodiments, a Chief Human Resources Officer (CHRO) digital twin 60638 (or HR digital twin 60638) is an executive digital twin configured for a human resources executive (e.g., a CHRO) of an enterprise or analogous executive tasked with overseeing the human resources HR aspects of the enterprise, such as a Chief People Officer (CPO), a chief talent officer, a head of human resources, a director of human resources, or the like. In example embodiments, the CHRO digital twin 60638 may depict different HR-related states of the enterprise, such as states relating to human capital management, workforce management, risk management, and the

management of payroll, recruitment, regulatory compliance, employee performance, benefits, employee relations, time and attendance, training and development, compensation, onboarding, offboarding, succession planning, and the like. In example embodiments, the CHRO digital twin 60638 may initially depict the various states at a lower granularity level. A user that is viewing the CHRO digital twin 60638 may select a state to drill down into the selected state and view the selected state at a higher level of granularity.

[2420] In example embodiments, the types of data that may be depicted in CHRO digital twin 60638 may include, but are not limited to: individual employee data, key performance indicators by business unit, key performance indicators by individual employee, risk management data, regulatory compliance data (e.g., OSHA and EPA compliance data), safety data, diversity data, benefits data (e.g., medical, dental, vision, and health savings accounts (HSA)) compensation data, compensation comparison data, compensation trend data, payroll data, overtime data, recruitment data, employee referrals data, applicant data, applicant screening data, applicant reference data, applicant background check data, offer data, time and attendance data, employee relations data, employee complaints data, onboarding data, offboarding data, employee training and development data, employee turnover rate data, voluntary employee turnover rate data, new hire turnover rate data, high performer turnover rate data, turnover rate by performance rating data, headcount and/or headcount planning data (e.g., headcount to plan percentage), promotion rate data, succession plan data, organizational levels data, span of control data, employee survey data, cost to move employees below midpoint data, comparative ratio data, simulation data, decision support data from AI and/or machine learning systems, prediction data from AI and/or machine learning systems, classification data from AI and/or machine learning systems, detection and/or identification data from AI and/or machine learning systems, and the like.

[2421] In example embodiments, a CHRO digital twin 60638 may depict a data item with an icon indicating whether the data item is at a normal state, a suboptimal state, a critical state, or an alarm state. In example embodiments, the icons may be different colors, fonts, symbols, codes or the like. For example, a CHRO digital twin 60638 may depict high performer turnover rate data with an orange icon indicating that the high performer turnover rate is at a critical level. Continuing the example, an HR executive may be enabled to escalate the high performer turnover rate data to another executive, such as the CEO, via the CHRO digital twin 60638. In example embodiments, a CHRO digital twin 60638 may automatically highlight data items that are at suboptimal, critical, or alarm state.

[2422] In example embodiments, a CHRO digital twin 60638 may be configured to provide an "in-twin" collaboration suite having tools that may facilitate communication and collaboration between enterprise stakeholders. In example embodiments, the "in-twin" collaboration tools may include an interface enabling a user to escalate and/or deescalate data sets to another user associated with the enterprise. In example embodiments, the interface may be configured to enable a user to send a message with the data set, generate a request or assign a task related to the data set, and/or schedule an event associated with the data set. In example embodiments, AI and/or machine learning could be leveraged to suggest message content, suggest event scheduling, suggest a

request or task, and/or suggest a request or task assignee. For example, an HR executive could escalate a data set related to employee training to the GC with a predictive text message about employee training and a calendar request at a time determined by AI and/or machine learning to attend a meeting related to employee training. In example embodiments, the "in twin" collaboration tools include digital twin conferences. In example embodiments, the "in twin" collaboration tools may include an "in-twin" messaging system and/or an "in-twin" video conferencing system for enabling enterprise stakeholders to communicate. In example embodiments, a machine learning and/or AI system may be leveraged for automatically generating and/or assigning tasks from these communications. In example embodiments, the "in-twin" videoconferencing system supports subchats. In example embodiments, the subchats may be created via a "drag-and-drop" action in the user interface. In example embodiments, the "in-twin" videoconferencing system may leverage machine learning and/or AI to make suggestions to optimize a user's lighting, audio, camera placement, and the like. In example embodiments, the "in twin" videoconferencing system leverages machine learning and/or AI to automatically disable the video feed upon the detection of an inappropriate activity in the video feed. In example embodiments, the "in twin" collaboration suite includes an "in-twin" stakeholder approval system for collecting approval on actions from other enterprise stakeholders. In example embodiments, "in-twin" collaboration tools may include an AI-driven translation system configured to intelligently translate communications amongst enterprise stakeholders to achieve maximum understanding by the user of the digital twin, wherein the AI driven translation system is configured to translate from a first language to a second language (e.g., translate English into a foreign language) and is also configured to translate terminology or jargon such that it is consumable by the user. These features described in connection with the CHRO digital twin 60638 may be deployed with other types of digital twins described herein, including ones for other executives, including to facilitate collaboration among different types of executives, such as for enterprise control tower activities, such as monitoring operations, development activities, or other aspects of the enterprise across locations, departments, and functions. Collaboration and communication tools and associated rules may be configured to use company-, industry and domain-specific taxonomies and lexicons when representing entities, states and flows within the CHRO digital twin 60638, such as ones relating to health and safety of workers, ones related to education and training, ones related to performance indicators, ones related to worker attributes (including psychographic, demographic and similar factors), and many others.

[2423] In example embodiments, a CHRO digital twin 60638 may be configured to identify, interview, select, hire, and onboard new employees. In some of these embodiments, the CHRO digital twin 60638 may be configured to research, track, and report on applicant data, including, but not limited to, employee referral data, applicant education data, applicant testing data, applicant experience data, applicant reference data, applicant screening data, applicant background check data, applicant interview data, job application data, applicant resume data, applicant cover letters, applicant offer data, and the like. The CHRO digital twin 60638 may interact with and share such applicant data and reporting with other executive digital twins, as described herein. The CHRO digital twin 60638 may include machine learning, AI, and/or other intelligence such as analytics,

to process job applications, resumes, cover letters, applicant reference materials, applicant screening data, applicant interview data, and the like in order to identify and select potential new employees and/or to identify other executives or enterprise stakeholders that may be interested in such information.

[2424] In example embodiments, the EMP 60000 may obtain HR-relevant data from the enterprise's human resources management software (e.g., via an API), human capital software, workforce management software, payroll software, applicant tracking software, accounting software, employee applicant software, publicly disclosed financial statements, third-party reports, tax filings, social media software, job listing websites, recruitment software, and the like.

[2425] In example embodiments, a CHRO digital twin 60638 may provide an interface for an HR executive to perform one or more HR-related workflows. For example, the CHRO digital twin 60638 may provide an interface for an HR executive to perform, supervise, or monitor workflows, the entities involved in the workflows, and attributes thereof, such as onboarding workflows, offboarding workflows, dismissal workflows, decision documentation workflows, succession planning workflows, candidate assessment workflows, candidate screening workflows, compliance workflows, disciplinary workflows, review workflows, interview workflows, offer workflows, employee training workflows, and many others.

[2426] In example embodiments, a CHRO digital twin 60638 may leverage an executive agent 60704 that is trained on a user's (e.g., an HR executive's) actions (e.g., behaviors, responses, interactions and preferences) using the expert agent system 60010 in response to events and situations encountered by the user (e.g., alerts, notifications, escalations, delegations, presentations of data, events, and the like). In some of these embodiments, the client application 60104 hosting the CHRO digital twin 60638 may report actions taken by the user in response to various events encountered by the user via the CHRO digital twin 60638. For example, the client application 60104 may identify events such as a request to authorize a new hire, a request to terminate an employee, or a notification indicating that employee turnover has reached a critical threshold. In this example, the client application 60104 may record and report the actions taken by the user in response to such events and may report the actions in relation to the identified events to the expert agent system 60010, as well as any other features that are relevant to the event. In response, the expert agent system 60010 may train an executive agent 60704 on behalf of the user, such that the executive agent may perform or recommend actions to the user when similar events are encountered in the future.

[2427] References to features and functions of the EMP and digital twins in this example of a human resources department and a CHRO digital twin 60638 should be understood to apply to other departments and digital twins, and their respective projects and workflows, except where context indicates otherwise.

[2428] In example embodiments, the executive digital twins may link to, interact with, integrate with and/or be used by a number of different applications. For example, the executive digital twins may be used in automated AI-reporting tools 60700, collaboration tools 60702, in connection with

executive agents 60704, in board meeting tools 60708, for training modules 60710, and for planning tools 60712.

[2429] In example embodiments, AI reporting tools 60700 assist users to report one or more states to another user. For example, a subordinate may need to report an identified issue to a higher-ranking member of the enterprise (e.g., CTO may wish to report an issue that needs to be addressed to the CEO). In example embodiments, the AI reporting tool 60700 may be configured to receive a request to report a state from a client device 60102. In example embodiments, the AI-reporting tool 60700 may identify the appropriate recipients of the reported state based on the type of request, the role of the user that issued the request and the organizational structure of the entity. In example embodiments, the AI-reporting tool may determine the role of the user and the recipients of the report from the organizational digital twin of the enterprise. In example embodiments, the AI-reporting tool 60700 may determine whether the intended recipients of a notification have access rights to the data being shared from the executive digital twin. For example, if the CFO is reporting to the CEO, it is likely that the CEO has access to all the enterprise's data and will not be precluded from receiving the report. Conversely, if the CFO wishes to delegate the handling of an issue via the AI-reporting tool to an employee in her business unit, the recipient may not have access to such data. In this scenario, the AI-reporting tool 60700 may notify the requesting user (e.g., the CFO) that certain types of data may not be shared with the subordinate employee and may determine a manner by which the issue may be reported to the subordinate without sharing the non-accessible data. Upon determining that a user has access rights to view a particular state of data, the AI-reporting tool 60700 may generate a report that is for the intended recipient. In example embodiments, the AI-reporting tool may leverage the NLP services of the intelligence system to generate the report. In example embodiments, the AI-reporting tool 60700 may leverage an executive agent 60704 to determine when to report a state and the appropriate recipients of the reported state. In these embodiments, the executive agent 60704 may be trained on interactions of the user with the client application 60104 and digital twins that were previously presented to the user.

[2430] In example embodiments, the AI-reporting tool 60700 may be configured to monitor one or more user-defined key performance indicators (KPIs). Examples of KPIs of an enterprise may include, but are not limited to, with respect to systems, facilities, processes, functions, or workforce units: uptime (e.g., of an assembly line or other manufacturing system), capacity utilization, on-standard operating efficiency, overall effectiveness, downtime, amount of unscheduled downtime, setup time, an amount of inventory turns, inventory accuracy, quality metrics relating to products and services, first-pass yield amounts for the enterprise, an amount of rework required, days-sales-outstanding (DSOs), an amount of scrap or waste produced, throughput, changeover, maintenance percentage, yield per system or unit, overall yield, industry reviews, industry ratings, customer reviews, customer ratings, editorial reviews, awards, social media and website attention metrics, search engine performance metrics, safety metrics, health metrics, environmental impact metrics, political metrics, certification and testing metrics, regulatory metrics, social impact metrics, financial and investment metrics, corporate bond ratings, trade association metrics, union metrics,

lobbying organization ratings, advertising performance metrics, referral metrics, and many others. Additional or alternative KPI metrics may be defined by a user. Examples of these KPI metrics may include an amount or percentage of failed audits, a number or percentage of deliveries that are on-time/late, a number of customer returns, a number of employee training hours, employee turnover percentage, number of reportable health or safety incidents, revenue per employee, profit per employee, schedule attainment metrics, total cycle time, and the like.

[2431] In example embodiments, the collaboration tools 60702 include various tools that allow collaboration between executives of the enterprise. In example embodiments, the collaboration tools include digital-twin enabled video conferencing. In these embodiments, the EMP 60000 may present participants in the video conference with the requested view of an enterprise digital twin. For example, during a Board meeting, a CTO proposing an update to the machinery or equipment in a factory may present an environment digital twin of the factory where the updates to the machinery or equipment would be made. In this example, the CTO may illustrate the results of simulations performed in the factory without the updates and with the updates. The simulation may illustrate how the update may benefit the enterprise using a number of selected metrics (e.g., throughput, profits, employee safety, or the like). Collaboration and communication tools and associated rules may be configured to use company-, industry- and domain-specific taxonomies and lexicons when representing entities, states and flows within the digital twin.

[2432] In example embodiments, executive agents 60704 are expert agents that are trained to perform tasks on behalf of executive users. As discussed, in example embodiments, a client application may monitor the user of the client application by a user when using the client application 60104. In these embodiments, the client application 60104 may monitor the states of an executive digital twin that the user drills down into, the states that the user reports to a superior and/or delegates to a team member in her respective business unit, decisions that are made, and the like. As the user uses the client application 60104, the expert agent system 60010 may train one or more machine-learned models on behalf of the particular user, such that the models may be leveraged by an executive agent 60704 to perform tasks on behalf of or recommend actions to the user.

[2433] In example embodiments, Board meeting tools 60708 are tools that are used to prepare for, to access within and/or to follow-up on board and similar meetings, such as Board of Directors, Board of Trustees, shareholder meetings, annual meetings, investor meetings, and other important meetings. References to Board meetings herein should be understood to encompass these and other important meetings that require executive preparation, attendance and/or attention. In example embodiments, Board meeting tools 60708 may allow different users to present one or more states of an enterprise digital twins within the context of a Board report or Board meeting. For example, a user (e.g., a COO) may share a simulation of a proposed logistics solution from the COO digital twin 60624 with one or more devices (e.g., a device in the Board room and/or devices of participants accessing the Board meeting remotely). In example embodiments, a Board meeting tool 60708 may limit access to certain types of data based on time, scope, and permissions. For example, a Board meeting tool 60708 may require that all geolocations that board members be

registered before a Board meeting (e.g., Board room, designated home offices for those joining by phone or video, and the like), such that some or all of the data depicted in a digital twin that is being presented can only be viewed on a device that is at one of the registered geolocations and/or only for a defined duration, such as from a few hours before through a few hours after a meeting, or only during the meeting. Similarly, in example embodiments, the Board meeting tools 60708 may limit access to some or all of the data shared in a presented digital twin to particular times (e.g., during the Board meeting or the day of the Board meeting). Other examples of board meeting tools 60708 are discussed throughout the application.

[2434] In example embodiments, training modules 60710 may include software tools that are used to train a user. In example embodiments, the training modules 60710 may leverage digital twins to improve executive training for an enterprise. For example, a training module 60710 may provide real-world examples that are based on the data collected from the enterprise. The training module 60710 may present the user with different scenarios via an executive digital twin 60710 and the user may take actions. Based on the actions, the training module 60710 may request a simulation from the EMP 60000, which in turn returns the results to the user. In this way, the user may be trained in scenarios that are based on the actual enterprise of the user.

[2435] In example embodiments, planning tools 60712 are software tools that leverage digital twins to assist users to make plans for the enterprise. In example embodiments, a planning tool 60712 may be configured to provide a graphical user interface that allows an executive to make plans (e.g., budgets, defining KPIs, etc.). In example embodiments, the planning tool 60712 may be configured to request a simulation from the IMP 60000 given the parameters set in the created plan. In response, the EMP 60000 may return the results of the simulation and the user can determine whether to adjust the plan. In this way, the user may iteratively refine the plan to achieve one or more objectives. In example embodiments, an executive agent 60704 may monitor and track the actions taken while the plan is being refined by the user so that the expert agent system 60010 may train the executive agent 60704 to generate or recommend plans to the user in the future.

[2436] The enterprise digital twins may be leveraged and/or interface with other software applications without departing from the scope of the disclosure.

[2437] Figure 382 illustrates an example implementation of the EMP 60000. In this example, the EMP 60000 is in communication with a plurality of client applications 60104 and a set of enterprise entities 60202. In the example, the EMP 60000 receives enterprise data from a set of enterprise entities 60202, such as a sensor system 60032, physical entities 60204, digital entities 60208, computational entities 60210, and/or network entities 60212 belonging to and/or associated with the enterprise. In example embodiments, the enterprise data may relate to environments, processes, and/or a condition of the enterprise. For example, a sensor system 60032 may be deployed within an enterprise factory whereby the sensor system 60032 provides sensor readings (e.g., vibration data, location data, motion data, temperature data, pressure data, or the like) relating to the factory in general or a piece of machinery, equipment, or other physical or workforce asset within the factory. Within the factory, a number of physical assets (e.g., robots, autonomous vehicles, smart equipment, personnel and the like) or other entities may output data streams relating to the

operation of the assets or other entities. Additionally or alternatively, the enterprise may include a number of digital assets (e.g., CRM, ERP, databases, or the like) that provide data streams relating to sales, costs, human resources or the like. The network entities may provide networking-related data, including bandwidth, API requests, throughput, detected cyber-attacks, or the like. The computational entities may provide data relating to a computing infrastructure of an enterprise. In example embodiments, the enterprise management system 60000 may receive data from other sources as well, including third-party data 60060 from third-party data providers. Taken in combination, the data from the enterprise entities 60202 and/or other data sources may provide information relating to the status of the industrial factory and the machinery contained therein, the state of various processes (e.g., industrial processes, sales workflows, hiring processes, logistics workflows, and the like), the efficiencies of the processes, the financial health of the enterprise, and the like.

[2438] In example embodiments, the enterprise entities may communicate directly with the EMP 60000 via a communication network. Additionally or alternatively, one or more of the enterprise assets may stream data to a local data collection system 60230 that collects and stores enterprise data locally. In example embodiments, the local data collection system 60230 may provide the collected data to an edge intelligence system 60232 of the enterprise.

[2439] In example embodiments, the edge intelligence system 60232 may be executed by an edge device 60064 configured to receive data, such as from the local data collection systems 60230, a local sensor system 60032, or other enterprise entities 60202 that are located in or near a physical location of the entities (e.g., at an industrial factory) and may perform one or more edge-related processes relating to the received data. The edge device may be a pre-configured and/or substantially self- or automatically configuring computing device, such as an "edge intelligence in a box" device. An edge-related process may refer to a process that is performed at an edge device in order to store sensor data, reduce bandwidth on a communication network, and/or reduce the computational resources required at a backend system. Examples of edge processes can include data filtering, signal filtering, data processing, compression, encoding, quick-predictions, quick-notifications, emergency alarming, and the like, and may include creation of automated smart data bands. For example, the edge intelligence system 60232 may determine whether to transmit a subset of the data to the EMP 60000 or to store the subset of the data locally until it is explicitly requested from the EMP 60000. In another example, the edge intelligence system 60232 may be configured to compress data streams (e.g., sensor data streams) to improve data throughput of high-volume data streams (e.g., vibration data). In example embodiments, the edge intelligence system 60232 may be configured to analyze the high-volume data to determine whether to compress or stream a raw data stream. In example embodiments, the local data collection system 60230 and the edge intelligence system 60232 may be embodied in edge devices 60064 of the enterprise. In example embodiments, the edge intelligence system 60232 may communicate data to the EMP 60000. In some of these embodiments, the edge intelligence system 60232 communicates data to the EMP 60000 via a network enhancement system 60234.

[2440] In example embodiments, the network enhancement system 60234 may be configured to optimize flow of data transmitted from one or both of the edge intelligence system 60232 and the local data collection system 60230 and received by the EMP 60000. For example, a local data collection system 60230 may be configured to collect data from one or more real world environments, entities, ecosystems, and/or processes, which may be analyzed by a connected edge intelligence system 60232. In this example, the edge intelligence system 60232 may transmit the collected data to the network enhancement system 60234, which may optimize transmission of the data to the EMP 60000 for processing and implementation by the EMP 60000. The EMP 60000 may store, analyze, or otherwise process the transmitted data to the client applications 60104, such that the client applications 60104 may update enterprise digital twins (e.g., role-based digital twins, environment digital twins, cohort digital twins, and the like) that are hosted by the client applications 60104.

[2441] In example embodiments, the network enhancement system 60234 may include one or more signal amplifiers, signal repeaters, digital filters, analog filters, digital-to-analog converters, analog-to-digital converter and/or antennae configured to optimize the flow of data. In example embodiments, the network enhancement system may include a wireless repeater system such as is disclosed by US Patent No. 7,623,826 to Pergal, the entirety of which is hereby incorporated by reference. The network enhancement system 60234 may optimize the flow of data by, for example, filtering data, repeating data transmission, amplifying data transmission, adjusting one or more sampling rates and/or transmission rates, and implementing one or more data communication protocols.

[2442] In example embodiments, the network enhancement system 60234 may include one or more processors configured to perform digital signal processing to optimize the flow of data. The one or more processors may implement optimization algorithms to optimize the flow of data. The one or more processors may determine one or more optimal paths in a network, the network enhancement system 60234 transmitting the data along the one or more optimal paths. The network enhancement system 60234 may be configured to implement a software filter via the one or more processors. The software filter may filter data before transmission to the EMP 60000, for example, to lower network bandwidth consumed by data transmission. The one or more processors may determine that portions of data are relevant only to one or more intended recipients, such as digital twins, executive agents, collaboration suites, or other components of the EMP 60000 and determine optimal paths based upon intended recipients of the portions of data.

[2443] In example embodiments, the network enhancement system 60234 may be configured to optimize data flow between a plurality of nodes over a plurality of data paths. In example embodiments, the network enhancement system 60234 may transmit a first portion of data over a first path of the plurality of data paths and a second portion of data over a second path of the plurality of data paths. The network enhancement system 60234 may determine that one or more data paths, such as the first data path, the second data path, other data paths, are advantageous for transmission of one or more portions of data. The network enhancement system 60234 may make determinations of advantageous data paths based upon one or more networking variables, such as

one or more types of data being transmitted, one or more protocols being suitable for transmission, present and/or anticipated network congestion, timing of data transmission, present and/or anticipated volumes of data being or to be transmitted, and the like. Protocols suitable for transmission may include transmission control protocol (TCP), user datagram protocol (UDP), and the like. In example embodiments, the network enhancement system may be configured to implement a method for data communication such as is disclosed by US Patent No. 9,979,664 to Ho et al., the entirety of which is hereby incorporated by reference.

[2444] The EMP 60000 receives enterprise data (e.g., directly or via the network enhancement system 60234, an edge intelligence system 60232, a local data collection system 60230 or from any other data source). In example embodiments, the digital twin system 60004 may structure and/or store the enterprise data in one or more digital twin databases (e.g., graph databases, relational databases, SQL databases, distributed databases, blockchains, caches, servers, and/or the like). In example embodiments, the client application 60104 requests an enterprise digital twin 60220 from the EMP 60000. In response, the digital twin system 60004 may generate and serve the requested enterprise digital twin 60220 (e.g., a role-based digital twin, executive digital twin, environment digital twin, process digital twin, cohort digital twins, or the like) to the client application 60104, whereby the enterprise digital twin 60220 may include the enterprise data and/or data that was derived from the enterprise data (e.g., by the intelligence services system). The client application 60104 may provide an interface for the user of the client application 60104 to interact with the requested digital twin 60220. For example, the user may delegate tasks relating to a depicted state to subordinates and/or may notify a superior of a depicted state via the digital twin interface. In another example, the user may drill down into a particular state and may initiate a corrective action via the digital twin interface. In example embodiments, the client application 60104 may allow the user to share the digital twin 60220 (or a portion thereof) within a collaboration tool 60224 or access collaboration features of a collaboration tool 60224 within the digital twin 60220. For example, the client application 60104 may allow the user to share a depicted state of the enterprise digital twin 60220 into a board meeting collaboration tool. Additionally or alternatively, an expert agent 60222 may monitor the interactions of the user with the digital twin and may report the interactions to the expert agent system 60010 of the EMP. In example embodiments, the expert agent system 60010 may receive the interactions and may train the expert agent 60222 based on the interactions with the digital twin, as well as outcomes stemming from the expert agent 60010. For example, the expert agent 60010 may be trained to identify situations where the user delegates tasks or notifies a superior.

[2445] The executive digital twins discussed with respect to FIG. 239 are provided for example and not intended to limit the scope of the disclosure. Additional and/or alternative data types may be included in a respective type of executive digital twin.

[2446] FIG. 240 illustrates an example method 60900 for configuring and serving an enterprise digital twin. In example embodiments, the method may be executed by the digital twin system 60004. The method may be performed with respect to different types of enterprise digital twins,

including role-based digital twins (e.g., executive digital twins), cohort digital twins, environment digital twins, process digital twins, and/or the like.

[2447] At 60902, the structural views for a particular type of digital twin are selected. In example embodiments, the structural views can be stored in a graph database (representing interconnected data) or in a geospatial database (representing coordinates of actual facilities).

[2448] At 60904, associated transactional data for the digital twin is selected. In example embodiments, a combination of interaction data and transaction data is selected at grain that is suitable for the dynamic interaction within the digital twin is selected. This selection process may involve dynamic configuration of the structure, functions and features of a data mart or other summarization system and/or may work dynamically using typically high-performance database storage mechanisms (such as columnar databases or in memory databases).

[2449] At 60908, embellishment and/or augmentation data for the digital twin is selected. In example embodiments, embellishment data are the associated attributes that can be tied to elements within the executive digital twin. For example, in generating an environment digital twin of a factory, embellishment or augmentation data may include the ages of machinery or other assets in the factory, the names of key third-party suppliers that could replace items with supply chain deliveries, the inputs or outputs of process flows that occur within the factory, identities of managers, indicators of states and flows, and many others. In an abstract executive digital twin, the embellishment data may include social media data, for example, sentiment analytics that can be associated with the customer hierarchical views.

[2450] At 60908, a representation medium for the digital twin is selected. In example embodiments, the final representation can be multi-faceted, this can include a range of devices from simple mobile phone-based devices and touchscreen tablets to special-purpose devices and/or immersive AR/VR headsets, among many others. The representation medium impacts the volume and nature of data that is preferably selected in the earlier steps. In example embodiments, selection of a representation medium is provided as a feedback indicator to the data and networking pipeline, such that filtering and data path selection can be undertaken with awareness of end device and other capabilities and requirements of the representation medium. This may occur automatically, such as by an agent that is trained to provide context-sensitive feedback based on a training set of outcomes.

[2451] At 60910, the perspective views are constructed. In example embodiments, the perspective builder 60312 generates a level and nature of data that allows for different types of user to interact with the digital twin while gaining the appropriate level of perspective. For example, with a CEO-level view the CEO may require the context of third-party alternatives, market forces, and current strategic initiatives. In this example, the perspective builder 60312 takes these considerations into account in producing the level of digital twin appropriate for the CEO, furthermore, this will impact the data selection process as different grains of data are appropriate for the different views. These different perspectives can be simultaneously interacted with various roles allowing the executive to provide their guidance on the same topic while seeing and interaction with information relevant to their specific needs.

[2452] At 60912, user notifications are enabled. In example embodiments, notifications within the digital twin are controlled by the grain of the data selected and the required perspective. For example, a CTO level view requires notifications of various technology changes and technology market forces, the CTO digital twin is constantly being overlaid with these notifications that are structurally associated with the relevant part of the digital environment abstract or concrete. In an organizational chart, for example, the CTO could be seeing the implementation options for new technology to provide more efficient communication between organizational units in a strategic planning exercise to acquire a new company. Simultaneously the CFO is seeing the financial impacts of these various options, and the CEO is being notified of decisions that might impact the future market opportunities regarding the upcoming company acquisition.

[2453] FIG. 241 illustrates example embodiments of a system 61000 for fault diagnosis in an industrial environment 61010 having several components 61012, 61014, 61016, ..., 61018. Herein, the term "industrial environment" may include industrial plant and built environments that may be closely connected with the running of an industrial enterprise, and may generally represent a facility for production of goods, independently or as part of a group of such facilities, and may for example involve an industrial process and chemical business, a manufacturing industry, food and beverage industry, agricultural industry, swimming pool industry, home automation industry, leather treatment industry, paper making process, and the like. For the purpose of the disclosure, the term "industrial environment" may be used to denote any process environment including the entire industrial plant or its subunits. Such industrial environment may usually include machinery which may employ individual or interconnected components 61012, 61014, 61016, ..., 61018. Many industrial processes and machines may be controlled and/or powered by electric motors. Such processes and machines may include pumps providing fluid transport for chemical and other processes, fans, conveyor systems, compressors, gear boxes, motion control devices, HVAC systems, screw pumps, and mixers, as well as hydraulic and pneumatic machines driven by motors. Such motors may be combined with other system components, such as valves, pumps, furnaces, heaters, chillers, conveyor rollers, fans, compressors, gearboxes, and the like, as well as with appropriate power control devices such as motor starters and motor drives, to form industrial machines and actuators. For example, an electric motor may be combined with a motor drive providing variable electrical power to the motor, as well as with a pump, whereby the motor may rotate the pump shaft to create a controllable pumping system.

[2454] All these components 61012, 61014, 61016, ..., 61018 may be prone to vibrations which affect the operation of the industrial environment 61010. These vibrations may arise due to various factors, such as, but not limited to, broken bearing in the motor, broken or cracked rotor bars in the motor, misalignment in the motor, imbalance in the motor, material build-up in the motor, etc. The system 61000 of the disclosure may be implemented for fault diagnosis due to, but not limited to, such factors. In example embodiments, rotor bar defects and weakening may be a precursor to secondary deterioration that may lead to further and costly repairs, such as replacement of a rotor core and the like. Therefore, by detecting broken or weakening rotor bars, maintenance and repair costs may be minimized. Knowing the count of rotor bars may be a factor in determining when

maintenance and/or service of one or more rotor bars may be best actioned. As an example, by applying a rotor bar failure rate to a formula that predicts when a rotor bar may fail, knowing a count of rotor bars for a given machine, among other things like cycle rate, age, and the like may facilitate predicting when conducting service and/or testing of rotor bar-based systems may beneficially be conducted. A predictive maintenance circuit may predict maintenance events for industrial and other machines and/or may predict maintenance for a machine with a greater number of rotor bars sooner than for a comparable machine with fewer rotor bars.

[2455] Further, as shown, the system 61000 may include sensors 61022, 61024, 61026, ..., 61028 associated to the industrial environment 61010. Each of the sensors 61022, 61024, 61026, ..., 61028 may be operatively coupled to at least one of the components 61012, 61014, 61016, ..., 61018. As shown, a 'sensor 1' 61022 may be operatively coupled to a 'component 1' 61012; a 'sensor 2' 61024 may be operatively coupled to a 'component 2' 61014; a 'sensor 3' 61026 may be operatively coupled to a 'component 3' 61016; ...; and a 'sensor n' 61028 may be operatively coupled to a 'component n' 61018. The system 61000 may include, connect to, or integrate with the sensors 61022, 61024, 61026, ..., 61028 that may connect to the components 61012, 61014, 61016, ..., 61018. In this manner, the sensors 61022, 61024, 61026, ..., 61028 may provide information about the industrial environment 61010, about one or more machines, components, or devices in the industrial environment 61010, about one or more network conditions (such as network bandwidth, spectrum availability, congestion, interference, cost, timing, and/or availability), or about one or more cloud conditions or parameters.

[2456] The sensors 61022, 61024, 61026, ..., 61028 may be configured to generate sensor data values in response to one or more sensed parameters. In example embodiments, the one or more sensed parameters may include at least one of a set of temperature parameters, pressure parameters, humidity parameters, wind parameters, rainfall parameters, tide parameters, storm surge parameters, cloud cover parameters, snowfall parameters, visibility parameters, radiation parameters, audio parameters, video parameters, image parameters, water level parameters, quantum parameters, flow rate parameters, signal power parameters, signal frequency parameters, motion parameters, velocity parameters, acceleration parameters, lighting level parameters, analyte concentration parameters, biological compound concentration parameters, metal concentration parameters, or organic compound concentration parameters. The effects of each of the listed parameters are described in the disclosure. In example embodiments, the sensors 61022, 61024, 61026, ..., 61028 may provide a stream of data over time that may have a phase component, such as relating to acceleration or vibration, allowing for the evaluation of phase or frequency analysis of different operational aspects of a piece or portion of equipment or an operating component. In example embodiments, the sensors 61022, 61024, 61026, ..., 61028 may provide a stream of data that may not be conventionally phase-based, such as temperature, humidity, load, and the like. In example embodiments, the sensors 61022, 61024, 61026, ..., 61028 may provide a continuous or near continuous stream of data over time, periodic readings, event-driven readings, and/or readings according to a selected interval or schedule.

[2457] In example embodiments, depending on the type of the components 61012, 61014, 61016, ..., 61018 being measured, the environment in which the components 61012, 61014, 61016, ..., 61018 are operating and the like, the sensors 61022, 61024, 61026, ..., 61028 may comprise one or more of, without limitation, a vibration sensor, an optical vibration sensor, a thermometer, a hygrometer, a voltage sensor, a current sensor, an accelerometer, a velocity detector, a light or electromagnetic sensor (e.g., determining temperature, composition and/or spectral analysis, and/or object position or movement), an image sensor, a structured light sensor, a laser-based image sensor, an infrared sensor, an acoustic wave sensor, a heat flux sensor, a displacement sensor, a turbidity meter, a viscosity meter, a load sensor, a tri-axial vibration sensor, an accelerometer, a tachometer, a fluid pressure meter, an air flow meter, a horsepower meter, a flow rate meter, a fluid particle detector, an acoustical sensor, a pH sensor, and the like, including, without limitation, any of the sensors described throughout this disclosure and the documents incorporated by reference. The sensors 61022, 61024, 61026, ..., 61028 may typically include at least a temperature sensor, a load sensor, a tri-axial sensor, and/or a tachometer.

[2458] In example embodiments, the sensors 61022, 61024, 61026, ..., 61028 for data collection in the industrial environment 61010 may include an analog crosspoint switch deployed in the industrial environment 61010, such as a turbine-based power system. Monitoring for vibration in turbine systems, such as hydro-power systems, has been demonstrated to provide advantages in reduction in down time. However, with a large number of areas to monitor for vibration, particularly for on-line vibration monitoring, including relative shaft vibration, bearings absolute vibration, turbine cover vibration, thrust bearing axial vibration, stator core vibrations, stator bar vibrations, stator end winding vibrations, and the like, it may be beneficial to select among this list over time, such as taking samples from sensors for each of these types of vibration a few at a time. The industrial environment 61010 that includes such analog crosspoint switch may provide this capability by connecting each vibration sensor to separate inputs of the analog crosspoint switch and configuring the switch to output a subset of its inputs. A vibration data processing system, such as a computer system, may determine which sensors to pass through the analog crosspoint switch and configure an algorithm to perform the vibration analysis accordingly. As an example, sensors for capturing turbine cover vibration may be selected in the analog crosspoint switch to be passed onto a system that may be configured with an algorithm to determine turbine cover vibration from the sensor signals. Upon completion of determining turbine cover vibration, the crosspoint switch may be configured to pass along thrust bearing axial vibration sensor signals and a corresponding vibration analysis algorithm may be applied to the data. In this way, each type of vibration may be analyzed by a single processing system that works cooperatively with an analog crosspoint switch to pass specific sensor signals for processing.

[2459] In example embodiments, the sensors 61022, 61024, 61026, ..., 61028 may include vibration sensors, such as an analog vibration sensor, a digital vibration sensor, a fixed digital vibration sensor, a tri-axial vibration sensor, a single axis vibration sensor, an optical vibration sensor, and/or a machine vision system. In other aspects, the sensors 61022, 61024, 61026, ..., 61028 may include one or more of touch ID, chemical, electrical, acoustic, vibration, acceleration,

velocity, position, light, motion, temperature, magnetic fields, gravity, humidity, moisture, pressure, electrical fields, and/or sound sensors. In particular example embodiments, the sensors 61022, 61024, 61026, ..., 61028 may include at least one vibration measurement sensor coupled to a motor of the corresponding component 61012, 61014, 61016, ..., 61018. In such an example, the one or more sensed parameters may include vibration parameters related to a wobble in the motor of the corresponding component 61012, 61014, 61016, ..., 61018. In example embodiments, based on the sensed parameters, the system 61000 may make a bearing life prediction, identify a bearing health parameter, identify a bearing performance parameter, determine a bearing health parameter (e.g., fault conditions), and the like. The system 61000 may identify wear on a bearing, identify a presence of foreign matter (e.g., particulates) in the bearings, identify air gaps or a loss of fluid in oil/fluid coated bearings, identify a loss of lubrication in a set of bearings, identify a loss of power for magnetic bearings, identify strain/stress of flexure bearings, and the like. The system 61000 may identify optimal operation parameters for a piece of equipment to extend bearing life. The system 61000 may further identify behavior (e.g., resonant wobble) at a selected operational frequency (e.g., shaft rotation rate) for the components 61012, 61014, 61016, ..., 61018.

[2460] In example embodiments, the system 61000 may access equipment specifications, equipment geometry, bearing specifications, bearing materials, anticipated state information for bearing types, operational history, historical detection values, and the like for use in assessing an output of the components 61012, 61014, 61016, ..., 61018. The system 61000 may buffer a subset of detection values, intermediate data such as time-based detection values transformed to frequency information, filtered detection values, identified frequencies of interest, and the like for a predetermined length of time. The system 61000 may periodically store certain detection values to enable tracking of component performance over time. In example embodiments, based on relevant operating conditions and/or failure modes that may occur as detection values approach one or more criteria, the system 61000 may store data based on the fit of data relative to one or more criteria, such as those described throughout this disclosure. Based on one sensor input meeting or approaching specified criteria or range, the system 61000 may store additional data such as revolutions per minute (RPM) data, component loads, temperatures, pressures, vibrations, and/or other sensor data of the types described throughout this disclosure.

[2461] Further, as shown in FIG. 241, the system 61000 may include a digital twin datastore 61100. Herein, the digital twin datastore 61100 may store digital twins of various industrial environments, and the objects, devices, sensors, and/or humans in the industrial environments. In particular, the digital twin datastore 61100 may include at least one industrial-environment digital twin 61110 corresponding to the industrial environment 61010 (as described in the disclosure). Further, the at least one industrial-environment digital twin 61110 may include component digital twins 61112, 61114, 61116, ..., 61118 corresponding to the components 61012, 61014, 61016, ..., 61018 in the industrial environment 61010. Specifically, each of the component digital twins 61112, 61114, 61116, ..., 61118 may correspond to one of the components 61012, 61014, 61016, ..., 61018 in the industrial environment 61010. As shown, a 'component digital twin 1' 61112 may correspond to the 'component 1' 61012 and the associated 'sensor 1' 61022; a 'component digital

twin 2' 61114 may correspond to the 'component 2' 61014 and the associated 'sensor 2' 61024; a 'component digital twin 3' 61116 may correspond to the 'component 3' 61016 and the associated 'sensor 3' 61026; ...; and a 'component digital twin n' 61118 may correspond to the 'component n' 61018 and the associated 'sensor n' 61028. In this manner, the at least one industrial-environment digital twin 61110 and the component digital twins 61112, 61114, 61116, ..., 61118 may provide information about the industrial environment 61010 and the components 61012, 61014, 61016, ..., 61018 in the industrial environment 61010, respectively, including about one or more machines, components, or devices in the industrial environment 61010, about one or more network conditions (such as network bandwidth, spectrum availability, congestion, interference, cost, timing, and/or availability), or about one or more cloud conditions or parameters.

[2462] In example embodiments, the component digital twins 61112, 61114, 61116, ..., 61118 may be digital representations of one or more real-world elements. The digital twins may be configured to mimic, copy, and/or model behaviors and responses of the real-world elements in response to inputs, outputs, and/or conditions of the surrounding or ambient environment. Data related to physical properties and responses of the real-world elements may be obtained, for example, via user input, sensor input, and/or physical modeling (e.g., thermodynamic models, electrodynamic models, mechanical models, etc.). Information for the digital twin may correspond to and be obtained from the one or more real-world elements corresponding to the digital twin. For example, in example embodiments, the digital twin may correspond to one real-world element that may be a fixed digital vibration sensor on a machine component, and vibration data for the digital twin may be obtained by polling or fetching vibration data measured by the fixed digital vibration sensor on the machine component. In a further example, the digital twin may correspond to real-world elements such that each of the elements may be a fixed digital vibration sensor on a machine component, and vibration data for the digital twin may be obtained by polling or fetching vibration data measured by each of the fixed digital vibration sensors on the real-world elements.

[2463] In example embodiments, the at least one industrial-environment digital twin 61110 and the component digital twins 61112, 61114, 61116, ..., 61118 may be generated from imported data from one or more data sources such that the imported data may correspond to an industrial environment. The industrial-environment digital twin 61110 representing the industrial environment 61010 may be generated based on the imported data. Further, the component digital twins 61112, 61114, 61116, ..., 61118 may be identified within the industrial environment 61010, and a set of discrete component digital twins 61112, 61114, 61116, ..., 61118 representing the components 61012, 61014, 61016, ..., 61018 within the industrial environment 61010 may be generated. The set of discrete component digital twins 61112, 61114, 61116, ..., 61118 may be embedded within the industrial-environment digital twin 61110. Information related to the components 61012, 61014, 61016, ..., 61018 for generating the corresponding component digital twins 61112, 61114, 61116, ..., 61118 (simulated elements) may be obtained, for example, by evaluating behavior of corresponding real-world elements components 61012, 61014, 61016, ..., 61018 using mathematical models or algorithms, from libraries that may define information and behavior of the components 61012, 61014, 61016, ..., 61018 (e.g., physics libraries, chemistry

libraries, or the like). In example embodiments, the component digital twins 61112, 61114, 61116, …, 61118 may be generated based on properties of the corresponding component imported from at least one of: respective manufacturers of the components, onboard libraries, crowdsourced material, or subscription marketplaces. That is, the required information about the properties of the corresponding component may be obtained from existing data sources. For example, herein, the manufacturers of the components may provide information about specification, engineering drawings, etc. of the component; the onboard libraries may include information about historical behavior of the component for its implementation in the industrial environment 61010; the crowdsourced material may include information obtained from various public data sources related to a specification, operations, engineering drawings, etc. of the component; and the subscription marketplaces may include information obtained from one or more non-public data sources, which may be specific to certain component types or the like.

[2464] In example embodiments, digital twins may be depicted in a number of different role-based view types. For example, a manufacturing facility device may be viewed in an "operator" view that depicts the facility in a manner that may be suitable for a facility operator, a "C-Suite" view that depicts the facility in a manner that is suitable for executive-level managers, a "marketing" view that depicts the facility in a manner that is suitable for workers in sales and/or marketing roles, a "board" view that depicts the facility in a manner that is suitable for members of a corporate board, a "regulatory" view that depicts the facility in a manner that is suitable for regulatory managers, and a "human resources" view that depicts the facility in a manner that is suitable for human resources personnel. In response to a request that indicates a view type, the system 61000 may be configured to retrieve data for each digital twin that corresponds to the view type.

[2465] In example embodiments, the system 61000 may further include an executive digital twin 61120 configured to provide forecasted financial information for a given component based, at least in part, on the system characteristics determined to be related to the recognized pattern. Various metrics of the industrial environment 61010 may be represented in the executive digital twin 61120, such as providing managers, maintenance workers, executives, inspectors, and others a visual indication of the overall risk of an unscheduled shutdown, as well as visual indicators of the components that may be at risk, or that may be contributing to increases in the probability of an unscheduled shutdown of a factory, plant, system, process, line, machine, workflow, or the like. This may allow managers and executives to drill down, obtain further information, and undertake actions that may reduce the risk. As one illustrative example, an executive may be presented with a view of a set of factories, with one factory being represented in the executive digital twin 61120 in a different color (such as bright red) based on that factory having a probability of unscheduled shutdown that may exceed a threshold (or simply that it may have the highest probability among a set of factories). This may direct the attention of the executive to that factory, thereby leading to further insight into operational choices that may have been missed if the executive were merely presented with raw data, a spreadsheet, or the like where the unscheduled shutdown probability may need to be calculated, inferred, or the like. In examples, the executive digital twin 61120 may help to highlight legal risks (such as safety violations or instances where status information about

operations may indicate a likelihood that the company may breach a contract (such as by failing to produce an output that may be required by a contract) and the like), inventory managers, procurement personnel, and the like; and for executives, such as CEOs, CTOs, COOs, CIOs, CDOs, CMOs, and the like, who may interact with the executive digital twin 61120 that may represent whole factories, or sets of factories, such as to identify risks and opportunities that may involve understanding interactions of elements and/or contributions of elements involving the industrial environment 61010 to overall operations of an enterprise, to its strategies, or the like.

[2466] In example embodiments, the system 61000 may further comprise an operator digital twin 61130 configured to provide workflow information for performing maintenance for the given component based, at least in part, on the system characteristics determined to be related to the recognized pattern. The operator digital twin 61130 may allow an operator (such as a factory manager) to be presented with the probabilities of unscheduled shutdown of various component machines and processes; for example, a pump that may be maintaining a vacuum of a critical semiconductor production process for the factory (or a biologics production process, or the like) may be identified as having a high risk of failure, such as based on vibration analysis that may indicate cavitation, in combination with other data sources, such as ones indicating the age of the pump and its maintenance and operating history. The pump may be highlighted in the operator digital twin 61130, such as in a view configured for the factory manager, such as by highlighting the pump in a bright color and by animating the pump with movement (e.g., shaking a visual element) that may indicate a vibration problem is the likely contributor to the risk of unscheduled shutdown of the pump (which may cascade to a failure of the vacuum, the failure of a critical production process, and the shutdown of the entire factory). As a result of attention being directed by the operator digital twin 61130 by visual cues (as compared to a spreadsheet or raw data output), the factory manager may direct (including by interacting with the pump in the digital twin, such as by touching it) attention to the pump for maintenance or replacement. An instruction or message provided by one user (such as the factory manager or executive) may result in a message, or highlighting, in a different digital twin or user interface or dashboard that may be configured for another user. For example, the pump, if flagged by the factory manager in a view of the factory, may appear in a service worker's digital twin 13734, such as showing a route to the pump and subsequently switching to a view that may guide the worker through inspection, maintenance, service, and/or replacement.

[2467] It may be appreciated that digital twins may be represented in a number of different forms. In example embodiments, the at least one industrial-environment digital twin 61110 and the component digital twins 61112, 61114, 61116, ..., 61118 may be visual digital twins that may be configured to be rendered in a visual manner. Herein, the at least one industrial-environment digital twin 61110 and the component digital twins 61112, 61114, 61116, ..., 61118 may be rendered by a computing device, such that a human user may view digital representations of the industrial environment 61010 and/or the components 61012, 61014, 61016, ..., 61018 in the industrial environment 61010, respectively. In example embodiments, the digital twin may be rendered and output to a display device. In some of these example embodiments, the digital twin may be rendered

in a graphical user interface, such that a user may interact with the digital twin. For example, a user may "drill down" on a particular element (e.g., a physical object or device) to view additional information regarding the element (e.g., a state of a physical object or device, properties of the physical object or device, or the like). In example embodiments, the digital twin may be rendered and output in a virtual reality display. For example, a user may view a three-dimensional (3D) rendering of an environment (e.g., using a monitor or a virtual reality headset). While doing so, the user may view/inspect digital twins of physical assets or devices in the environment. In example embodiments, a data structure of the visual digital twins (e.g., digital twins that may be configured to be displayed in a two-dimensional (2D) or 3D manner) may include surfaces (e.g., splines, meshes, polygon meshes, or the like). In example embodiments, the surfaces may include texture data, shading information, and/or reflection data. In this way, a surface may be displayed in a more realistic manner. In example embodiments, such surfaces may be rendered by a visualization engine when the digital twin is within a field of view and/or when existing in a larger digital twin (e.g., a digital twin of an industrial environment). In example embodiments, a digital twin (and any digital twins embedded therein) may be represented in a non-visual representation (or "data representation"), as described elsewhere in this disclosure.

[2468] The system 61000 may further include one or more processors 61200 to implement the processes described herein, that may be deployed in part or in whole through a machine that may execute computer software, program codes, and/or instructions on the processor 61200. In example embodiments, the processor 61200 may be part of a server, cloud server, client, network infrastructure, mobile computing platform, stationary computing platform, or other computing platform. The processor 61200 may be any kind of computational or processing device capable of executing program instructions, codes, binary instructions, and the like. The processor 61200 may be or may include a signal processor, digital processor, embedded processor, microprocessor, or any variant such as a co-processor (math co-processor, graphic co-processor, communication co-processor, and the like) and the like that may directly or indirectly facilitate execution of program code or program instructions stored thereon. In example embodiments, a non-transitory computer readable storage medium may be provided having instructions stored thereon which, when executed across one or more processors 61200, causes at least a portion of the one or more processors to perform operations as described herein. In addition, the processor 61200 may enable execution of multiple programs, threads, and codes. The threads may be executed simultaneously to enhance performance of the processor and to facilitate simultaneous operations of the application.

[2469] The system may further include a client application 61300 which may be installed on a client device. In example embodiments, the client application 61300 may be configured to provide a digital representation and/or visualization of the digital twins as described herein. In example embodiments, the client application 61300 may include one or more software modules that may be executed by one or more server devices. These software modules may be configured to quantify properties of the digital twin, model properties of a digital twin, and/or to visualize digital twin behaviors. In example embodiments, these software modules may enable a user to select a

particular digital twin behavior visualization for viewing. In example embodiments, these software modules may enable a user to select to view a digital twin behavior visualization playback. In example embodiments, the client application 61300 may provide a selected behavior visualization of the digital twin. In example embodiments, the system 61000 may update the properties of the digital twin and/or one or more embedded digital twins on behalf of the client application 61300. In example embodiments, the client application 61300 may be an application relating to an industrial component or environment (e.g., monitoring an industrial facility or a component within, simulating an industrial environment, or the like). In example embodiments, the disclosure includes outputting visual digital twins to the client application 61300 that may display the visual digital twins via a virtual reality headset. In example embodiments, the disclosure includes outputting the visual digital twins to the client application 61300 that may display the visual digital twins via a display of a client device. In example embodiments, the disclosure includes outputting the visual digital twins to the client application 61300 that may display the visual digital twins via an augmented reality-enabled device.

[2470] FIG. 242 illustrates example embodiments of a computer-implemented process 62000 for fault diagnosis in an industrial environment (such as, the industrial environment 61000) having components (such as, the components 61012, 61014, 61016, ..., 61018). At 62002, several sensors (such as, the sensors 61022, 61024, 61026, ..., 61028) may be provided to the industrial environment 61000. Each of the sensors 61022, 61024, 61026, ..., 61028 may be operatively coupled to at least one of the components 61012, 61014, 61016, ..., 61018 and may be configured to generate sensor data values in response to one or more sensed parameters.

[2471] At 62004, the sensor data values may be processed to determine a recognized pattern therefrom. The processing may include the sensing performance value that may further include a determination of one or more of the following: a signal-to-noise performance for detecting a value of interest in the industrial system; a network utilization of the sensors in the industrial system; an effective sensing resolution for a value of interest in the industrial system; a power consumption value for a sensing system in the industrial system where the sensing system may include the sensors; a calculation efficiency for determining a secondary value; an accuracy and/or a precision of the secondary value; a redundancy capacity for determining the secondary value; and/or a lead time value for determining the secondary value. Example and non-limiting calculation efficiency values may include one or more determinations such as: processor operations to determine the secondary value; memory utilization for determining the secondary value; a number of sensor inputs from the number of sensors for determining the secondary value; and/or supporting data long term storage for supporting the secondary value. The recognized pattern may further be determined by performing an operation such as: determining a signal effectiveness of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to a value of interest; determining a sensitivity of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive confidence of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive delay time of at least one sensor of

the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive accuracy of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; determining a predictive precision of at least one sensor of the sensed parameter group and the updated sensed parameter group relative to the value of interest; and/or updating the recognized pattern value in response to external feedback. Examples and non-limiting values of interest may include: a virtual sensor output value; a process prediction value; a process state value; a component prediction value; a component state value; and/or a model output value having sensor data values as an input.

[2472] In example embodiments, the process 62000 may also include determining if the recognized pattern may relate to a system characteristic including at least one of: a fault operation for a given component of the components, an off-nominal operation for the given component of the components, or an exceedance value for the given component of the components. In such an example, the process 62000 may further include generating a notification in the client application 61300 in response to the determination that the recognized pattern may relate to the system characteristic for the given component. Example notifications may include: a notification to an operator, a notification to a user, a notification to a portable device associated with a user, a notification to a node of a network, a notification to a cloud computing device, a notification to a plant computing device, and/or a provision of the alert as external data to an offset system. Example and non-limiting notification conditions may include: a component of the system operating in a fault condition, a process of the system operating in a fault condition, a commencement of the utilization of cache storage and/or intermediate storage for sensor values due to a network communication limit, a change in the sensor data transmission protocol (including changes of a selected type), and/or a change in the sensor data transmission protocol that may result in loss of data fidelity or resolution (e.g., compression of data, condensing of data, and/or summarizing data).

[2473] At 62006, the at least one industrial-environment digital twin 61110 corresponding to the industrial environment 61010 along with the component digital twins 61112, 61114, 61116, ..., 61118 (where each of the component digital twins corresponds to one of the components in the industrial environment) may be retrieved. The one or more digital twins may be retrieved, including any embedded digital twins, from the digital twin datastore 61100. In example embodiments, retrieving the one or more digital twins may include identifying one or more dynamic models based on one or more properties that may be depicted in digital twins indicated by a request and a respective type of the one or more digital twins.

[2474] At 62008, the at least one industrial-environment digital twin 61110 and at least one respective component digital twin of the component digital twins 61112, 61114, 61116, ..., 61118 may be updated based on the sensor data values, at least in part, in response to a determination of the recognized pattern for the corresponding component. Herein, the one or more sensors from the sensors 61022, 61024, 61026, ..., 61028 may be selected which may provide necessary sensor data values for each corresponding associated component from the components 61012, 61014, 61016, ..., 61018, for updating of the corresponding component digital twins 61112, 61114, 61116, ..., 61118. For the sake of reducing the processing load, such sensor data values may be fetched only

in response to, e.g., when it has been determined that the corresponding component may be manifesting at least one of the recognized patterns. In example embodiments, the recognized pattern may include one or more of: broken bearing in the motor, broken or cracked rotor bars in the motor, misalignment in the motor, imbalance in the motor, material build-up in the motor, and/or the like. Such recognized patterns may indicate some kind of fault condition in the associated component with the corresponding sensor values; and therefore, it may be required to update the corresponding digital twin (one or more of the component digital twins 61112, 61114, 61116, ..., 61118), and additionally, the overall industrial-environment digital twin 61110. It may be appreciated that such an update may enable a digital representation of an industrial entity and/or environment where the real time digital representation may be a visualization of the digital twin.

[2475] In this example, the digital twin datastore 61100 may retrieve the digital twin of the machine and any embedded digital twins, such as any embedded motor digital twins and bearing digital twins, and any digital twins that embed the machine digital twin, such as the manufacturing facility digital twin. Further, in some examples, the dynamic model input data sources (e.g., sensor data and any other suitable data) may be selected based on available data sources (e.g., available sensors from the sensors 61022, 61024, 61026, ..., 61028) and the one or more required inputs of the dynamic model(s). In the present example, the retrieved dynamic model(s) may take one or more vibration sensor measurements from vibration sensors as inputs to the dynamic models. Further, the system 61000 may retrieve one or more measurements from each of the selected data sources and may run the dynamic model(s) using the retrieved vibration sensor measurements as inputs, and may calculate one or more outputs that may represent bearing vibration fault level states. Next, the system 61000 may update one or more bearing fault level states of the component digital twins 61112, 61114, 61116, ..., 61118 as well as the at least one industrial-environment digital twin 61110 based on the one or more outputs of the dynamic model(s).

[2476] At 62010, a request may be received from the client application 61300 to check an operational condition of a particular component from the components in the industrial environment 61010. In example embodiments, the request may be received from the client application 61300 that corresponds to the industrial environment 61010 and/or one or more industrial entities within the industrial environment 61010. In example embodiments, the system 61000 may receive requests from the client application 61300 to update properties of a digital twin in order to enable a digital representation of an industrial entity and/or environment where the real time digital representation may be a visualization of the digital twin. For this purpose, based on the request from the client application 61300, the one or more digital twins may be required to fulfill the request (e.g., as determined and retrieved from the digital twin datastore 61100). The client application 61300 may obtain vibration fault level states of the bearings and may display the obtained vibration fault level state associated with each bearing and/or display colors associated with fault level severity (e.g., red for alarm, orange for critical, yellow for suboptimal, green for normal operation) in the rendering of one or more of the digital twins on a display interface.

[2477] At 62012, the at least one industrial-environment digital twin 61110 and the at least one respective component digital twin 61112, 61114, 61116, ..., 61118 corresponding to the particular

component in the client application 61300 may be rendered in response to the received request and based on the operational condition of the particular component. For this purpose, based on the request from the client application 61300, the one or more digital twins may be required to fulfill the request (e.g., as determined and retrieved the from digital twin datastore 61100, and the same may be rendered). In example embodiments, the process 62000 may further include configuring the client application 61300 to allow for selection of the notification. Herein, the rendering the at least one industrial-environment digital twin 61110 and the at least one respective component digital twin 61112, 61114, 61116, ..., 61118 corresponding to the given component is in response to the selection of the notification. In an example, the notification may be in the form of a text message with a link to be clicked (selected), a notification message to be clicked, a heads-up notification to be clicked, a pop-up notification to be clicked, a voice notification to be responded in a voice command, and/or the like.

[2478] In example embodiments, the rendering may further include executing a simulation for the at least one industrial-environment digital twin and the at least one respective component digital twin based on the recognized pattern. This provides that whenever the recognized pattern indicates a fault condition (as described in the disclosure), the updated digital twin may be simulated to depict the fault condition for perusal of the client. In further example embodiments, the rendering may further include executing another second simulation for the at least one industrial-environment digital twin and the at least one respective component digital twin based on a normal operation of the corresponding component. This provides that along with simulation of the fault condition of the component, another simulation may be shown side-by-side which may depict the normal operation of the corresponding component; and thus the client may be able to clearly distinguish an effect of the fault condition in the operation of the given component. Herein, the simulation may simulate an effect of the recognized pattern on an operation of the corresponding component. In example embodiments, a digital twin may be rendered by a computing device, such that a human user may view the digital representations of real-world industrial assets, devices, workers, processes, and/or environments. For example, the digital twin may be rendered and outcome to a display device. In example embodiments, dynamic model outputs and/or related data may be overlaid on the rendering of the digital twin. In example embodiments, dynamic model outputs and/or related information may appear with the rendering of the digital twin in a display interface. In example embodiments, the related information may include real-time video footage associated with the real-world entity represented by the digital twin. In example embodiments, the related information may include a sum of each of the vibration fault level states in the machine. In example embodiments, the related information may be graphical information. In example embodiments, the graphical information may depict motion and/or motion as a function of frequency for individual machine components. In example embodiments, graphical information may depict motion and/or motion as a function of frequency for individual machine components, where a user may be enabled to select a view of the graphical information in the x, y, and z dimensions. In example embodiments, graphical information may depict motion and/or motion as a function of frequency for individual machine components, where the graphical information may include harmonic peaks

and peaks. In example embodiments, the related information may be cost data, including the cost of downtime per day data, cost of repair data, cost of new part data, cost of new machine data, and the like. In example embodiments, related information may be a probability of downtime data, probability of failure data, and the like. In example embodiments, related information may be time to failure data. In example embodiments, the related information may be recommendations and/or insights. For example, recommendations or insights received from the cognitive intelligence system related to a machine may appear with the rendering of the digital twin of a machine in a display interface.

[2479] In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin to the client application may be via a display device of a user device. In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin to the client application may be via an augmented reality-enabled device. In example embodiments, the rendering the at least one industrial-environment digital twin and the at least one respective component digital twin to the client application may be via a virtual reality headset. That is, as described in the disclosure, digital twins may be rendered in a number of different forms. In example embodiments, the digital twin may be rendered and output to a display device. In some of these example embodiments, the digital twin may be rendered in a graphical user interface, such that a user may interact with the digital twin. For example, a user may "drill down" on a particular element (e.g., a physical object or device) to view additional information regarding the element (e.g., a state of a physical object or device, properties of the physical object or device, or the like). In example embodiments, the digital twin may be rendered and output in a virtual reality display. For example, a user may view a 3D rendering of an environment (e.g., using a monitor or a virtual reality headset). While doing so, the user may view/inspect digital twins of physical assets or devices in the environment. In example embodiments, the client application 40070 may be an augmented reality application, whereby the requested digital twin may be depicted in an AR-enabled device. In these example embodiments, the requested digital twin may be filtered such that visual elements and/or text may be overlaid on the display of the AR-enabled device.

[2480] As discussed in example embodiments, the disclosure includes a process for fault diagnosis in an industrial environment having a plurality of components including providing a plurality of sensors to the industrial environment where each of the plurality of sensors may be operatively coupled to at least one of the plurality of components and configured to generate a plurality of sensor data values in response to one or more sensed parameters; processing the plurality of sensor data values to determine a recognized pattern therefrom; retrieving at least one industrial-environment digital twin corresponding to the industrial environment where the at least one industrial-environment digital twin may include a plurality of component digital twins, with each of the plurality of component digital twins corresponding to one of the plurality of components in the industrial environment, and where the at least one industrial-environment digital twin and the plurality of component digital twins may be visual digital twins that may be configured to be rendered in a visual manner; updating the at least one industrial-environment digital twin and at

least one respective component digital twin of the plurality of component digital twins based on the plurality of sensor data values, at least in part, in response to a determination of the recognized pattern for the corresponding component; receiving a request from a client application to check an operational condition of a particular component from the plurality of components in the industrial environment; and rendering the at least one industrial-environment digital twin and the at least one respective component digital twin corresponding to the particular component in the client application in response to the received request and based on the operational condition of the particular component.

[2481] In example embodiments, the request may be received from a client application that may correspond to an industrial environment and/or one or more industrial entities within the industrial environment. In example embodiments, the request may be received from a client application that supports an Industrial Internet of Things (IIoT) sensor system. In example embodiments, the digital twins may be digital twins of at least one of industrial entities and industrial environments. In example embodiments, the dynamic models may take data selected from the set of vibration, temperature, pressure, humidity, wind, rainfall, tide, storm surge, cloud cover, snowfall, visibility, radiation, audio, video, image, water level, quantum, flow rate, signal power, signal frequency, motion, displacement, velocity, acceleration, lighting level, financial, cost, stock market, news, social media, revenue, worker, maintenance, productivity, asset performance, worker performance, worker response time, analyte concentration, biological compound concentration, metal concentration, and/or organic compound concentration data.

[2482] Referring to FIGS. 243 to 248, in combination, an implementation of the system 61000 and the process 62000 is provided. Referring to FIG. 243, a schematic 53000 is shown to depict a vibration and issue detection scheme for an electric motor (such as an electric motor 63100, 63200 as shown in FIGS. 244 and 245) using principles described in the disclosure. For example, FIG. 243 shows specifically a phase (e.g., a failing phase) based on the vibration and issue detection scheme or process for the electric motor. In particular, the system may watch for "bad poles" usually due to physical impact on the windings. The wobble induced by the bad poles may be relatively very slight and may not be detected by pedestrian electric motor testers (as there may be minimal or no interruption in continuity), but the system may detect the wobble (even so slight) and determine bad (or failing or even less efficient) poles based on the signal detected and ruling out other items. As shown in FIG. 246, a client/user 63300 may take a picture or video (or other media recording) of a component (e.g., electric motor 63100, 63200) using a client device 63310. As used herein, the term client/user 63300 may generally apply to any entity utilizing the data processing system described herein, such as a person (e.g., an individual) interacting with an application program or an application program itself, for example, performing automated tasks. The client 63300 may access the system via the client device 63310 (e.g., a personal computer, tablet, or smartphone) to view information about and/or manage a digital twin in accordance with any of the example embodiments described herein. Also, as used herein, client devices 63310, including those associated with the system and any other device described herein, may exchange information via any communication network which may be one or more of a Local Area Network

("LAN"), a Metropolitan Area Network ("MAN"), a Wide Area Network ("WAN"), a proprietary network, a Public Switched Telephone Network ("PSTN"), a Wireless Application Protocol ("WAP") network, a BLUETOOTH™ network, a wireless LAN network, and/or an Internet Protocol ("IP") network such as the Internet, an intranet, or an extranet. Note that any devices described herein may communicate via one or more such communication networks. In example embodiments, an interactive graphical display interface may let the client 63300 define and/or adjust certain parameters and/or provide or receive automatically generated recommendations or results. While the following description may often refer to a graphical user interface (GUI) intended to present information to and receive information from a person, it should be understood that in many cases, the same functionality may be provided through a non-graphical user interface, such as a command line and, further, similar information may be exchanged with a non-person user via a programming interface.

[2483] FIGS. 247 and 248 provide GUIs 63400 and 63500, respectively, which display a digital twin of the component selected by the client device. In example embodiments, the client application may be configured to provide a digital representation and/or visualization of the digital twin of an industrial entity. In example embodiments, the client application may include one or more software modules that may be executed by one or more server devices. These software modules may be configured to quantify properties of the digital twin, model properties of the digital twin, and/or to visualize digital twin behaviors. In example embodiments, these software modules may enable a user to select a particular digital twin behavior visualization for viewing. In example embodiments, these software modules may enable a user to select to view a digital twin behavior visualization playback. In example embodiments, the client application may provide a selected behavior visualization of the digital twin. Interfaces and dashboards may include graphical interfaces (such as for laptops, tablets, and mobile devices), touch screen interfaces, voice-activated interfaces, augmented reality interfaces, virtual reality interfaces, mixed reality interfaces, application programming interfaces (APIs), and the like. Digital twins may be of various types, such as component digital twins represent an individual part of component; machine digital twins that represent an entire machine; system digital twins that represent a system involving multiple components, parts, machines or the like and their interactions; worker digital twins that represent one or more attributes or states of a set of workers; arrangement digital twins that represent the layout or arrangement of entities (such as, without limitation, the arrangement of components, assets, machines, workers, or other elements on a factory floor); augmented, virtual and/or mixed reality digital twins that provide a realistic experience for a user, such as simulating or mimicking interaction with an asset, another worker, a workflow, or the like (such as for training a worker or group of workers how to operate or undertake maintenance on a machine or system, how to undertake a workflow involving a machine or system, or the like); abstract digital twins (such as ones that represent elements and relationships, such as in topologies, hierarchies, flow diagrams, or the like), and others. In example embodiments, interfaces and dashboards may be provided that facilitate drilling down and/or zooming up in a digital twin (whether under user control or by automation, such as based on an understanding of status information, contextual

information, user interactions, or other factors), such as to obtain a relatively more detailed view of a component of a larger view (e.g., to see a specific part of a machine in an exploded view); to move up to a wider view that encompasses more components and/or their interactions; to obtain additional information (such as to view additional metrics related to a metric represented in a digital twin, more granular data, source data that may be used to determine a metric, or the like); and the like. In example embodiments, interfaces and dashboard may be configured to facilitate switching between views or types of digital twin of the same entities (whether under user control or by automation, such as based on an understanding of status information, contextual information, user interactions, or other factors involving the digital twin). For example, a user may switch from an overall schematic view that may represent current status information for the machines and workflows on a factory floor to a 3D view that shows a realistic representation of one of the machines (such as one that may have been highlighted as having an issue, such as where a data collector has determined that it may be operating outside normal parameters for temperature, vibration, pressure, or the like).

[2484] FIG. 249 illustrates a unified architecture 64000 for implementation of the processes as described herein. Further, FIG. 250 illustrates a data collection architecture 64100 for the unified architecture 64000, FIG. 251 illustrates an artificial intelligence (AI) architecture 64200 for the unified architecture 64000, and FIG. 252 illustrates a data storage and display architecture 64300 for the unified architecture 64000. The following description has been explained in reference to FIGS. 249 to 252 in combination.

[2485] As illustrated, the architecture 64100 may include sensors 64110 having a 1D vibration probe, a 2D vibration probe, a 3D vibration probe, and other sensors. Further, a data interface 64120 may be provided to transmit the sensor data generated by the sensors 64110 for processing. As used herein, the data interface 64120 may refer to the number of distinct buses and the type of data buses used by the memory, to transfer data during read and write operations. The type of characteristic of the data bus may refer to whether the data bus is uni-directional or bi-directional. The two data bus configurations commonly implemented in synchronous random-access memories (RAM) (e.g., static random-access memories (SRAMs)) are common input/output (I/O) and separate I/O. The common I/O data bus configuration may include one bi-directional data bus used to transfer data for both read and write operations. The separate I/O data bus configuration may include two uni-directional data buses, one used to transfer data for read operations and one used to transfer data for write operations. The architecture 64100 may also include machine learning models 64130 which may be housed relatively very close to the edge, potential even with the sensors 64110, and apply machine learning at the edge. Using edge computing, more particularly, hybrid centralized server may be used to implement an intelligent system for anomaly detection in a process environment using an edge computing model that has been found to provide some relatively interesting benefits and efficiencies. In the edge, network and cloud, self-organization of network protocol selection, data storage, routing of collection workloads, generation of alerts, augmentation of interfaces, and many other features may result in enhanced performance and reliability across the network and for the digital twin itself. Intelligent networking may automate

control of packet rates based on channel conditions to avoid problems of interference and free riding in heavy industrial environments without requiring complicated RF hardware filters. These features and more may result in improved throughput, latency, power efficiency and reliability for data-intensive requirements of cloud-based AI systems and digital twins. Further, as shown, a multiplexer 64140 may be provided for selective transmission of data, for example, that may switch between vibration sensors dynamically to manage overall data flow volume. The architecture 64100 may also include an edge signal processor 64150 which may process transformations such as time to frequency transformations at the edge to reduce overall dataflows. Also, a network protocol selection layer 64160 may be provided which may execute a base transportation that may be applied to allow for the movement of data. In example embodiments, the architecture 64100 may further include a mesh network bandpass 64170 to provide mesh network and filtering repeaters for a fault tolerant high-performance communications platform. Further, the architecture 64100 may include a storage 64180 to store data streamed in a large-scale data storage location, in which the data may be stored as a data lake or data warehouse. This stored data may become the existing data and may be applied to build new models.

[2486] As illustrated, the architecture 64200 may include new data 64210 streaming in from the sensors 64110 which may be applied to the existing models and existing data 64220 which may be applied to build or train new models. The architecture 64200 may also include a central signal processing unit 64230 which may apply techniques such as Fourier transform, Hilbert transform, Cepstrum, or envelope spectrum techniques to filter and sort the raw input data, with at least one key component being the transformation from time space into frequency space. Herein, the signal processed data may be used for further model processing. The architecture 64200 may further include a feature extraction layer 64240 which may combine other features such as power and temperature into input variables. These combinations may include statistical analysis or processing methods, for example principal component analysis or variable standardization. The architecture 64200 may further include a pattern recognition layer 64250 which may provide machine learning engines that may be trained on existing data to produce models that provide human interpretable and actionable outputs. Herein, new data streams may feed through to produce new machine learning generated models. Thereby, the architecture 64200 may generate unsupervised models 64252 which may be machine learning models where the predicted variable may not be used as part of the training process, but oftentimes may be used for classification; and supervised models 64254 which may be machine learning models where the predicted variable may be used as part of the training process.

[2487] As illustrated, the architecture 64300 may include a data source 64310 which may, in turn, include multiple data sources like 'Data Source A', 'Data Source B', 'Data Source C' and the like. These data sources may come from a variety of other areas that may be relevant to modeling and digital twin display mechanism. The architecture 64300 may also include maps 64320 which may provide 2D and 3D maps, which in particular are 2D and 3D models that may form a spatial or abstract structure for the digital twins. The architecture 64300 may further provide digital twin tools 64330 which may provide digital twin display and interaction functionality (e.g., using a

digital twin display and interaction engine), and may hold a digital twin of the real-world objects overlaid with data, including vibration data. The architecture 64300 may also provide classification (e.g., using a classification engine 64340) which may be a machine learning engine that may classify input data into buckets, as may be effectively achieved by unsupervised models. The architecture 64300 may also provide a fault engine 64350 which may be a machine learning engine that produces a prediction model showing likely faults, and which may be displayed in the digital twin or may be used as an alerting mechanism.

[2488] Such AI-driven digital twins may enable actionable insights about a wide range of industrial systems and processes involving many types, and combinations, of sensor data. Sensor data generated, monitored, analyzed, and enhanced by AI may be used for intelligent flagging, alerts, and guidance in digital twins, as well as in other presentation layers. This in turn, may optimize scheduling and routing of expert workers and components required for service and operational tasks, thus improving uptime, and reducing cost while mitigating the growing constraints caused by a shrinking workforce of industrial experts. Field service experts may be guided by artificial intelligence, such as for data collection, tagging and characterization, and undertaking of maintenance and other service tasks, mitigating problems of a shrinking workforce of industrial experts. Real-time visualization of conditions of machines, components and environments may also be combined with financial information about assets to provide executive level decision support. At the operator level, the digital twin may provide real-time guidance of operations, maintenance, and repairs, as well as in-field service activities. For manufacturers, product development, sales, and marketing, activities may be enhanced through digital twins that may supply aggregate or fleet-based information about real-world performance and machine conditions. Key features may allow for streamlined design, development, and implementation of AI-driven digital twins without disrupting installed IT systems. The underlying platform may bring together rich inputs for artificial intelligence through flexible multiplexing of different combinations of sensor channels, while avoiding problems of overload in data and network channels. These sensor networks may be deployed in a mesh network for coordinated data collection among multiple data collection units. Once connected, automated band selection based on context and expert-based rules may enable multi-tenant applications, such as involving IoT devices from different vendors, multi-provider wireless networking, and the like, where resources may be securely allocated between private and public access uses. The digital twin platform may readily be extended to deploy AI systems that may replicate human industrial expertise for monitoring, forecasting, reporting, guidance, and control, including populating augmentations to digital twins.

[2489] Further, the shared data collection and handling architecture may enable customization of digital twin interfaces, workflows and applications for specific objects, workflows, and roles. For example, machine digital twins may provide real-time visualization of conditions of machines, components, and environments, with configuration and augmentation driven and configured by artificial intelligence operating on rich industrial data sources. Such understanding of machine behavior, as determined by analytic algorithms operating on vibration data streams may be translated into a 3D animation of the diagnosed machine behavior, which may then be integrated

into a digital twin for display. For example, a user of the digital twin may observe a prompt (such as a diagnostic result indicating a health problem with a machine) and may then observe the vibration pattern of the machine, along with its base motion, by clicking through to the embedded animation. Other digital twins may include executive digital twins that may be configured for asset owners and senior executives that may include versions that represent real-time and forecasted financial information about assets and decision support based on AI-expert processing of sensor data and other inputs. Operator digital twins may be configured for operators like floor managers to provide real time representation of sensed machine conditions and AI-augmented workflow guidance for operations, maintenance, and repairs. Service digital twins may be configured for field service workers to provide rich detail about machine characteristics, conditions, and workflow guidance for field service activities, including data collection, tagging, and service work. Manufacturer digital twins may be configured for original equipment manufacturers to provide aggregate or fleet-based information about real-world performance and machine condition and guidance for product development, sales, and marketing activities. Access controls, interfaces, and other elements may be controlled at all levels of the data handling architecture, allowing the benefits of shared data collection, storage, and processing while maintaining security and control.

[2490] In one practical implementation, data acquisition may amass large raw continuous digital streams of multi-axis vibration data on each machine bearing. Analytic scans may be performed on the data streams for granular identification of the stream windows and frequency regions exhibiting high frictional and impact energy. Custom filtering techniques may further refine the data, leaving the frictional and non-random impacting elements. Each data stream may be mined with diverse processing and filtering techniques for the earliest fault detection possible, so that simple preventative actions, like lubrication, may be used to defuse potentially adverse, or even terminal, machine outcomes. This and other similar approaches may provide early identification of bearing lubrication issues that may have resulted from frictional forces, which may be separated from impact modulations that predominantly manifest from bearing defects. Soon after starting these techniques in the field, rising friction levels may be detected on pump bearings (e.g., two pump bearings) located in the same area of a chemical plant, prompting a visual inspection. Oil levels may be relatively very low, an issue that may likely be missed during a last preventive check. A modest frictional increase may trigger sensitive analytics and alert the digital twin operator to the problem. In another practical implementation, increased friction may be detected, and although grease lubrication was adequate, it may be discovered that the outer race of the bearing may be spinning within the housing, generating the previously undetected friction.

[2491] In real-world examples for detecting vibrations in a roller bearing, a process may involve filtering out interference noise which may mask the bearing impact signals, which may be further subdivided into random and periodic varieties. The interference noise that may potentially mask the signal(s) may consist predominantly of running speed and electrical line frequency harmonics. After filtering, only the frictional and synchronous impacting forces may remain and may be clearly identified (and also visualized) as emanating from the bearing. The next step may be to scan across frequencies as well as across a span of the stream for the high frequency band that may

provide a best response. For example, picking a frequency band between about 1 and about 20 kHz that may be approximately about 500-1 kHz wide where there may be the best signal-to-noise ratio for impacting vibration. This may replace historical methodology of manually pre-selecting an arbitrary band or bands based on generalized machine structural characteristics or a specific accelerometer's resonant frequency (e.g., the IMI 629M05 triaxial accelerometer). It may be appreciated that bearing impacts may cause modulations of multiple carriers in the waveform. The most commonly observed carrier utilized may be the bearing housing's ringing frequency, which may be usually unique for every machine. Generally, the carrier may tend to be lower for bigger machines and higher for smaller machines (e.g., inversely proportional to the mass) but as a practical matter, it may not be possible to know with assurance. A favorable approach may be to locate a carrier frequency region with a relatively best signal-to-noise response for impact energy after interference signals have been removed. The next step may include removing harmonics of machines and electrical frequencies. Herein, all harmonics of the running speed as well as any electrical frequencies may be removed prior to a scan. To achieve this for higher frequency bands, it may be important that the running speeds and electrical frequencies be known with acceptable precision or ascertainable from the factory environment, connections between machines, and the like. For electrical frequencies, precision may be typically required for components such as variable-speed drives, as line frequencies may be generally known. A degree of precision may be needed because the confidence of determining higher-order harmonics may be a function of the precision of the fundamental frequency component. In these examples, sufficiently high sampling resolution and data stream length (optionally reinforced with a center-of-moment spectral bin distribution calculation where applicable) may be utilized to determine these fundamentals as accurately as possible. The next step may include performing high-resolution envelope analysis to extract modulations within the filtered band. Afterwards, algorithms may be employed for extracting a modulation index (e.g., ratio of modulating signal(s) to carrier frequency) and additional filtering may be performed in order to effectively differentiate between the largely random impacts associated with bearing lubrication issues, versus relatively more regular periodic impacts that occur at non-integer running speed harmonic frequencies, typically associated with bearing defects such as spalls on the inner and outer races, bearing cage defects, and the like. It may be appreciated that improper lubrication may generate random impacting noise at markedly increased and/or erratic levels that allows for detection of unusual behavior that may be outside that of the typical machine and its harmonic patterns. At this point, the analysis may reveal elevated frictional forces within the bearing itself. Periodic impacting (bumping) may occur at a certain rate, and this means it may be modulating periodically; that is, non-random bumping suggests a bearing defect. In contrast, frictional forces may be different in that they occur largely randomly with respect to time, rather than regularly. The conveyance of this information regarding the bearing defect energy may be especially useful through the employment of a digital color-coded animation of the bearing movement at slow motion bearing impact frequencies which may be either periodic or random in nature. The animation, along with a diagnosis and a prompt for inspection, may be displayed in a digital twin of the machine.

[2492] With the system 61000, it may be possible to consider multiple components 61012, 61014, 61016, …, 61018 in the industrial environment 61010, each having its own micro-characteristics and not only average measures of the components as traditionally implemented. Moreover, it may be possible to accurately monitor and continually assess the health of individual components, predict their remaining functional lives, and even estimate a health and a remaining useful functional life of each of the components 61012, 61014, 61016, …, 61018 in the industrial environment 61010. This may be a significant advance for applied prognostics; and discovering a system and methodology to do so in an accurate and efficient manner may help reduce unplanned down time for complex systems (resulting in cost savings and increased operational efficiency). It may also be possible to achieve optimal control of the industrial environment 61010 if the functional life of the components 61012, 61014, 61016, …, 61018 may be accurately determined.

[2493] The simulation in the system 61000 may be configured to process simulation models corresponding to the digital twins using multiprocessor computer systems. These simulation models may be implemented, for example, using a Bayesian filtering framework. In example embodiments, a simulation platform may be configured to execute each respective simulation model using simulation engines executing in parallel across processors on the multiprocessor computer system. In example embodiments, the system may further include a data platform which may be configured to process data query tasks using the multiprocessor computer systems. This data platform may utilize techniques such as a map-reduce programming to process each of the data query tasks.

[2494] Referring to FIG. 253, the artificial intelligence system 65050 may define a machine learning model 65052 for performing analytics, simulation, decision making, and prediction making related to data processing, data analysis, simulation creation, and simulation analysis of one or more of the manufacturing entities 65010. The machine learning model 65052 is an algorithm and/or statistical model that performs specific tasks without using explicit instructions, relying instead on patterns and inference. The machine learning model 65052 builds one or more mathematical models based on training data to make predictions and/or decisions without being explicitly programmed to perform the specific tasks. The machine learning model 65052 may receive inputs of sensor data as training data, including event data 65140 and state data 65140 related to one or more of the manufacturing entities 65010. The sensor data input to the machine learning model 65052 may be used to train the machine learning model 65052 to perform the analytics, simulation, decision making, and prediction making relating to the data processing, data analysis, simulation creation, and simulation analysis of the one or more of the manufacturing entities 65010. The machine learning model 65052 may also use input data from a user or users of the information technology system. The machine learning model 65052 may include an artificial neural network, a decision tree, a support vector machine, a Bayesian network, a genetic algorithm, any other suitable form of machine learning model, or a combination thereof. The machine learning model 65052 may be configured to LEARN through supervised learning, unsupervised learning, reinforcement learning, self learning, feature learning, sparse dictionary learning, anomaly detection, association rules, a combination thereof, or any other suitable algorithm for learning.

[2495] The artificial intelligence system 65050 may also define the digital twin system 65070 to create a digital replica of one or more of the manufacturing entities 65010. The digital twin system 65070, the artificial intelligence system 65050, and the adaptive edge intelligence system 65060 can be included in the adaptive intelligence system 65080. The adaptive intelligence system 65080 can connect to the manufacturing entities 65010 through connectivity facilities 65020, which also permits connectivity with a monitoring system 65100 and a data collector system 65110. The digital replica of the one or more of the manufacturing entities may use substantially real-time sensor data to provide for substantially real-time virtual representation of the manufacturing entity and provides for simulation of one or more possible future states of the one or more manufacturing entities. The digital replica exists simultaneously with the one or more manufacturing entities 65010 being replicated. The digital replica provides one or more simulations of both physical elements and properties of the one or more manufacturing entities being replicated and the dynamics thereof, in example embodiments, throughout the lifestyle of the one or more manufacturing entities being replicated. The digital replica may provide a hypothetical simulation of the one or more manufacturing entities, for example during a design phase before the one or more manufacturing entities are constructed or fabricated, or during or after construction or fabrication of the one or more manufacturing entities by allowing for hypothetical extrapolation of sensor data to simulate a state of the one or more manufacturing entities, such as during high stress, after a period of time has passed during which component wear may be an issue, during maximum throughput operation, after one or more hypothetical or planned improvements have been made to the one or more manufacturing entities, or any other suitable hypothetical situation. In example embodiments, the machine learning model 65052 may automatically predict hypothetical situations for simulation with the digital replica, such as by predicting possible improvements to the one or more manufacturing entities, predicting when one or more components of the one or more manufacturing entities may fail, and/or suggesting possible improvements to the one or more manufacturing entities, such as changes to timing settings, arrangement, components, or any other suitable change to the manufacturing entities. The digital replica allows for simulation of the one or more manufacturing entities during both design and operation phases of the one or more manufacturing entities, as well as simulation of hypothetical operation conditions and configurations of the one or more manufacturing entities. The digital replica allows for invaluable analysis and simulation of the one or more manufacturing entities, by facilitating observation and measurement of nearly any type of metric, including temperature, wear, light, vibration, etc. not only in, on, and around each component of the one or more manufacturing entities, but in some embodiments within the one or more manufacturing entities. In example embodiments, the machine learning model 65052 may process the sensor data including the event data 65140 and the state data 65130 from a data storage system 65120 to define simulation data for use by the digital twin system 65070. The machine learning model 65052 may, for example, receive state data 65130 and event data 65140 related to a particular manufacturing entity of the plurality of manufacturing entities and perform a series of operations on the state data 65130 and the event data 65140 to format the state data 65140 and the event data 65140 into a format suitable for use by the digital

twin system 65070 in creation of a digital replica of the manufacturing entity. For example, one or more manufacturing entities may include a robot configured to augment products on an adjacent assembly line. The machine learning model 65052 may collect data from one or more sensors positioned on, near, in, and/or around the robot. The machine learning model 65052 may perform operations on the sensor data to process the sensor data into simulation data and output the simulation data to the digital twin system 65070. The digital twin system 65070 may use the simulation data to create one or more digital replicas of the robot, the simulation including for example metrics including temperature, wear, speed, rotation, and vibration of the robot and components thereof. The simulation may be a substantially real-time simulation, allowing for a human user of the information technology to view the simulation of the robot, metrics related thereto, and metrics related to components thereof, in substantially real time. The simulation may be a predictive or hypothetical situation, allowing for a human user of the information technology to view a predictive or hypothetical simulation of the robot, metrics related thereto, and metrics related to components thereof.

[2496] In example embodiments, the machine learning model 65052 and the digital twin system 65070 may process sensor data and create a digital replica of a set of manufacturing entities of the plurality of manufacturing entities to facilitate design, real-time simulation, predictive simulation, and/or hypothetical simulation of a related group of manufacturing entities. The digital replica of the set of manufacturing entities may use substantially real-time sensor data to provide for substantially real-time virtual representation of the set of manufacturing entities and provide for simulation of one or more possible future states of the set of manufacturing entities. The digital replica exists simultaneously with the set of manufacturing entities being replicated. The digital replica provides one or more simulations of both physical elements and properties of the set of manufacturing entities being replicated and the dynamics thereof, in embodiments throughout the lifestyle of the set of manufacturing entities being replicated. The one or more simulations may include a visual simulation, such as a wire-frame virtual representation of the one or more manufacturing entities that may be viewable on a monitor, using an augmented reality (AR) apparatus, or using a virtual reality (VR) apparatus. The visual simulation may be able to be manipulated by a human user of the information technology system, such as zooming or highlighting components of the simulation and/or providing an exploded view of the one or more manufacturing entities. The digital replica may provide a hypothetical simulation of the set of manufacturing entities, for example during a design phase before the one or more manufacturing entities are constructed or fabricated, or during or after construction or fabrication of the one or more manufacturing entities by allowing for hypothetical extrapolation of sensor data to simulate a state of the set of manufacturing entities, such as during high stress, after a period of time has passed during which component wear may be an issue, during maximum throughput operation, after one or more hypothetical or planned improvements have been made to the set of manufacturing entities, or any other suitable hypothetical situation. In example embodiments, the machine learning model 65052 may automatically predict hypothetical situations for simulation with the digital replica, such as by predicting possible improvements to the set of manufacturing

entities, predicting when one or more components of the set of manufacturing entities may fail, and/or suggesting possible improvements to the set of manufacturing entities, such as changes to timing settings, arrangement, components, or any other suitable change to the manufacturing entities. The digital replica allows for simulation of the set of manufacturing entities during both design and operation phases of the set of manufacturing entities, as well as simulation of hypothetical operation conditions and configurations of the set of manufacturing entities. The digital replica allows for invaluable analysis and simulation of the one or more manufacturing entities, by facilitating observation and measurement of nearly any type of metric, including temperature, wear, light, vibration, etc. not only in, on, and around each component of the set of manufacturing entities, but in some embodiments within the set of manufacturing entities. In example embodiments, the machine learning model 65052 may process the sensor data including the event data 65140 and the state data 65140 to define simulation data for use by the digital twin system 65070. The machine learning model 65052 may, for example, receive state data 65130 and event data 65140 related to a particular manufacturing entity of the plurality of manufacturing entities and perform a series of operations on the state data 65130 and the event data 65140 to format the state data 65140 and the event data 65140 into a format suitable for use by the digital twin system 65070 in the creation of a digital replica of the set of manufacturing entities. For example, a set of manufacturing entities may include a die machine configured to place products on a conveyor belt, the conveyor belt on which the die machine is configured to place the products, and a plurality of robots configured to add parts to the products as they move along the assembly line. The machine learning model 65052 may collect data from one or more sensors positioned on, near, in, and/or around each of the die machines, the conveyor belt, and the plurality of robots. The machine learning model 65052 may perform operations on the sensor data to process the sensor data into simulation data and output the simulation data to the digital twin system 65070. The digital twin system 65070 may use the simulation data to create one or more digital replicas of the die machine, the conveyor belt, and the plurality of robots, the simulation including for example metrics including temperature, wear, speed, rotation, and vibration of the die machine, the conveyor belt, and the plurality of robots and components thereof. The simulation may be a substantially real-time simulation, allowing for a human user of the information technology to view the simulation of the die machine, the conveyor belt, and the plurality of robots, metrics related thereto, and metrics related to components thereof, in substantially real time. The simulation may be a predictive or hypothetical situation, allowing for a human user of the information technology to view a predictive or hypothetical simulation of the die machine, the conveyor belt, and the plurality of robots, metrics related thereto, and metrics related to components thereof.

[2497] In example embodiments, the machine learning model 65052 may prioritize collection of sensor data for use in digital replica simulations of one or more of the manufacturing entities. The machine learning model 65052 may use sensor data and user inputs to train, thereby learning which types of sensor data are most effective for creation of digital replicate simulations of one or more of the manufacturing entities. For example, the machine learning model 65052 may find that a particular manufacturing entity has dynamic properties such as component wear and throughput

affected by temperature, humidity, and load. The machine learning model 65052 may, through machine learning, prioritize collection of sensor data related to temperature, humidity, and load, and may prioritize processing sensor data of the prioritized type into simulation data for output to the digital twin system 65070. In example embodiments, the machine learning model 65052 may suggest to a user of the information technology system that more and/or different sensors of the prioritized type be implemented in the information technology near and around the manufacturing entity being simulation such that more and/or better data of the prioritized type may be used in simulation of the manufacturing entity via the digital replica thereof.

[2498] In example embodiments, the machine learning model 65052 may be configured to LEARN to determine which types of sensor data are to be processed into simulation data for transmission to the digital twin system 65070 based on one or both of a modeling goal and a quality or type of sensor data. A modeling goal may be an objective set by a user of the information technology system or may be predicted or learned by the machine learning model 65052. Examples of modeling goals include creating a digital replica capable of showing dynamics of throughput on an assembly line, which may include collection, simulation, and modeling of, e.g., thermal, electrical power, component wear, and other metrics of a conveyor belt, an assembly machine, one or more products, and other components of the manufacturing ecosystem. The machine learning model 65052 may be configured to learn to determine which types of sensor data are necessary to be processed into simulation data for transmission to the digital twin system 65070 to achieve such a model. In example embodiments, the machine learning model 65052 may analyze which types of sensor data are being collected, the quality and quantity of the sensor data being collected, and what the sensor data being collected represents, and may make decisions, predictions, analyses, and/or determinations related to which types of sensor data are and/or are not relevant to achieving the modeling goal and may make decisions, predictions, analyses, and/or determinations to prioritize, improve, and/or achieve the quality and quantity of sensor data being processed into simulation data for use by the digital twin system 65070 in achieving the modeling goal.

[2499] In example embodiments, a user of the information technology system may input a modeling goal into the machine learning model 65052. The machine learning model 65052 may LEARN to analyze training data to output suggestions to the user of the information technology system regarding which types of sensor data are most relevant to achieving the modeling goal, such as one or more types of sensors positioned in, on, or near a manufacturing entity or a plurality of manufacturing entities that is relevant to the achievement of the modeling goal is and/or are not sufficient for achieving the modeling goal, and how a different configuration of the types of sensors, such as by adding, removing, or repositioning sensors, may better facilitate achievement of the modeling goal by the machine learning model 65052 and the digital twin system 65070. In example embodiments, the machine learning model 65052 may automatically increase or decrease collection rates, processing, storage, sampling rates, bandwidth allocation, bitrates, and other attributes of sensor data collection to achieve or better achieve the modeling goal. In example embodiments, the machine learning model 65052 may make suggestions or predictions to a user of the information technology system related to increasing or decreasing collection rates,

processing, storage, sampling rates, bandwidth allocation, bitrates, and other attributes of sensor data collection to achieve or better achieve the modeling goal. In example embodiments, the machine learning model 65052 may use sensor data, simulation data, previous, current, and/or future digital replica simulations of one or more manufacturing entities of the plurality of manufacturing entities to automatically create and/or propose modeling goals. In example embodiments, modeling goals automatically created by the machine learning model 65052 may be automatically implemented by the machine learning model 65052. In example embodiments, modeling goals automatically created by the machine learning model 65052 may be proposed to a user of the information technology system, and implemented only after acceptance and/or partial acceptance by the user, such as after modifications are made to the proposed modeling goal by the user.

[2500] In example embodiments, the user may input the one or more modeling goals, for example, by inputting one or more modeling commands to the information technology system. The one or more modeling commands may include, for example, a command for the machine learning model 65052 and the digital twin system 65070 to create a digital replica simulation of one manufacturing entity or a set of manufacturing entities, may include a command for the digital replica simulation to be one or more of a real-time simulation, and a hypothetical simulation. The modeling command may also include, for example, parameters for what types of sensor data should be used, sampling rates for the sensor data, and other parameters for the sensor data used in the one or more digital replica simulations. In example embodiments, the machine learning model 65052 may be configured to predict modeling commands, such as by using previous modeling commands as training data. The machine learning model 65052 may propose predicted modeling commands to a user of the information technology system, for example, to facilitate simulation of one or more of the manufacturing entities that may be useful for the management of the manufacturing entities and/or to allow the user to easily identify potential issues with or possible improvements to the manufacturing entities.

[2501] In example embodiments, the machine learning model 65052 may be configured to evaluate a set of hypothetical simulations of one or more of the manufacturing entities. The set of hypothetical simulations may be created by the machine learning model 65052 and the digital twin system 65070 as a result of one or more modeling commands, as a result of one or more modeling goals, one or more modeling commands, by prediction by the machine learning model 65052, or a combination thereof. The machine learning model 65052 may evaluate the set of hypothetical simulations based on one or more metrics defined by the user, one or more metrics defined by the machine learning model 65052, or a combination thereof. In example embodiments, the machine learning model 65052 may evaluate each of the hypothetical simulations of the set of hypothetical simulations independently of one another. In example embodiments, the machine learning model 65052 may evaluate one or more of the hypothetical simulations of the set of hypothetical simulations in relation to one another, for example by ranking the hypothetical simulations or creating tiers of the hypothetical simulations based on one or more metrics.

[2502] In example embodiments, the machine learning model 65052 may include one or more model interpretability systems to facilitate human understanding of outputs of the machine learning model 65052, as well as information and insight related to cognition and processes of the machine learning model 65052, i.e., the one or more model interpretability systems allow for human understanding of not only "what" the machine learning model 65052 is outputting, but also "why" the machine learning model 65052 is outputting the outputs thereof, and what process led to the machine learning model 65052 formulating the outputs. The one or more model interpretability systems may also be used by a human user to improve and guide training of the machine learning model 65052, to help debug the machine learning model 65052, to help recognize bias in the machine learning model 65052. The one or more model interpretability systems may include one or more of linear regression, logistic regression, a generalized linear model (GLM), a generalized additive model (GAM), a decision tree, a decision rule, RuleFit, Naive Bayes Classifier, a K-nearest neighbors algorithm, a partial dependence plot, individual conditional expectation (ICE), an accumulated local effects (ALE) plot, feature interaction, permutation feature importance, a global surrogate model, a local surrogate (LIME) model, scoped rules, i.e. anchors, Shapley values, Shapley additive explanations (SHAP), feature visualization, network dissection, or any other suitable machine learning interpretability implementation. In example embodiments, the one or more model interpretability systems may include a model dataset visualization system. The model dataset visualization system is configured to automatically provide to a human user of the information technology system visual analysis related to distribution of values of the sensor data, the simulation data, and data nodes of the machine learning model 65052.

[2503] In example embodiments, the machine learning model 65052 may include and/or implement an embedded model interpretability system, such as a Bayesian case model (BCM) or glass box. The Bayesian case model uses Bayesian case-based reasoning, prototype classification, and clustering to facilitate human understanding of data such as the sensor data, the simulation data, and data nodes of the machine learning model 65052. In example embodiments, the model interpretability system may include and/or implement a glass box interpretability method, such as a Gaussian process, to facilitate human understanding of data such as the sensor data, the simulation data, and data nodes of the machine learning model 65052.

[2504] In example embodiments, the machine learning model 65052 may include and/or implement testing with concept activation vectors (TCAV). The TCAV allows the machine learning model 65052 to LEARN human-interpretable concepts, such as "running," "not running," "powered," "not powered," "robot," "human," "truck," or "ship" from examples by a process including defining the concept, determining concept activation vectors, and calculating directional derivatives. By learning human-interpretable concepts, objects, states, etc., TCAV may allow the machine learning model 65052 to output useful information related to the manufacturing entities and data collected therefrom in a format that is readily understood by a human user of the information technology system.

[2505] In example embodiments, the machine learning model 65052 may be and/or include an artificial neural network, e.g., a connectionist system configured to "LEARN" to perform tasks by

considering examples and without being explicitly programmed with task-specific rules. The machine learning model 65052 may be based on a collection of connected units and/or nodes that may act like artificial neurons that may in some ways emulate neurons in a biological brain. The units and/or nodes may each have one or more connections to other units and/or nodes. The units and/or nodes may be configured to transmit information, e.g., one or more signals, to other units and/or nodes, process signals received from other units and/or nodes, and forward processed signals to other units and/or nodes. One or more of the units and/or nodes and connections therebetween may have one or more numerical "weights" assigned. The assigned weights may be configured to facilitate learning, i.e., training, of the machine learning model 65052. The weights assigned weights may increase and/or decrease one or more signals between one or more units and/or nodes, and in some embodiments may have one or more thresholds associated with one or more of the weights. The one or more thresholds may be configured such that a signal is only sent between one or more units and/or nodes, if a signal and/or aggregate signal crosses the threshold. In example embodiments, the units and/or nodes may be assigned to a plurality of layers, each of the layers having one, or both, of inputs and outputs. A first layer may be configured to receive training data, transform at least a portion of the training data, and transmit signals related to the training data and transformation thereof to a second layer. A final layer may be configured to output an estimate, conclusion, product, or other consequence of processing of one or more inputs by the machine learning model 65052. Each of the layers may perform one or more types of transformations, and one or more signals may pass through one or more of the layers one or more times. In example embodiments, the machine learning model 65052 may employ deep learning and being at least partially modeled and/or configured as a deep neural network, a deep belief network, a recurrent neural network, and/or a convolutional neural network, such as by being configured to include one or more hidden layers.

[2506] In example embodiments, the machine learning model 65052 may be and/or include a decision tree, e.g., a tree-based predictive model configured to identify one or more observations and determine one or more conclusions based on an input. The observations may be modeled as one or more "branches" of the decision tree, and the conclusions may be modeled as one or more "leaves" of the decision tree. In example embodiments, the decision tree may be a classification tree. the classification tree may include one or more leaves representing one or more class labels, and one or more branches representing one or more conjunctions of features configured to lead to the class labels. In example embodiments, the decision tree may be a regression tree. The regression tree may be configured such that one or more target variables may take continuous values.

[2507] In example embodiments, the machine learning model 65052 may be and/or include a support vector machine, e.g., a set of related supervised learning methods configured for use in one or both of classification and regression-based modeling of data. The support vector machine may be configured to predict whether A NEW example falls into one or more categories, the one or more categories being configured during training of the support vector machine.

[2508] In example embodiments, the machine learning model 65052 may be configured to perform regression analysis to determine and/or estimate a relationship between one or more inputs and one

or more features of the one or more inputs. Regression analysis may include linear regression, wherein the machine learning model 65052 may calculate a single line to best fit input data according to one or more mathematical criteria.

[2509] In example embodiments, inputs to the machine learning model 65052 (such as a regression model, Bayesian network, supervised model, or other type of model) may be tested, such as by using a set of testing data that is independent from the data set used for the creation and/or training of the machine learning model, such as to test the impact of various inputs to the accuracy of the model 65052. For example, inputs to the regression model may be removed, including single inputs, pairs of inputs, triplets, and the like, to determine whether the absence of inputs creates a material degradation of the success of the model 65052. This may assist with recognition of inputs that are in fact correlated (e.g., are linear combinations of the same underlying data), that are overlapping, or the like. Comparison of model success may help select among alternative input data sets that provide similar information, such as to identify the inputs (among several similar ones) that generate the least "noise" in the model, that provide the most impact on model effectiveness for the lowest cost, or the like. Thus, input variation and testing of the impact of input variation on model effectiveness may be used to prune or enhance model performance for any of the machine learning systems described throughout this disclosure.

[2510] In example embodiments, the machine learning model 65052 may be and/or include a Bayesian network. The Bayesian network may be a probabilistic graphical model configured to represent a set of random variables and conditional independence of the set of random variables. The Bayesian network may be configured to represent the random variables and conditional independence via a directed acyclic graph. The Bayesian network may include one or both of a dynamic Bayesian network and an influence diagram.

[2511] In example embodiments, the machine learning model 65052 may be defined via supervised learning, i.e., one or more algorithms configured to build a mathematical model of a set of training data containing one or more inputs and desired outputs. The training data may consist of a set of training examples, each of the training examples having one or more inputs and desired outputs, i.e., a supervisory signal. Each of the training examples may be represented in the machine learning model 65052 by an array and/or a vector, i.e., a feature vector. The training data may be represented in the machine learning model 65052 by a matrix. The machine learning model 65052 may learn one or more functions via iterative optimization of an objective function, thereby learning to predict an output associated with new inputs. Once optimized, the objective function may provide the machine learning model 65052 with the ability to accurately determine an output for inputs other than inputs included in the training data. In example embodiments, the machine learning model 65052 may be defined via one or more supervised learning algorithms such as active learning, statistical classification, regression analysis, and similarity learning. Active learning may include interactively querying, by the machine learning model 65052, a user and/or an information source to label new data points with desired outputs. Statistical classification may include identifying, by the machine learning model 65052, to which a set of subcategories, i.e., subpopulations, A NEW observation belongs based on a training set of data containing

observations having known categories. Regression analysis may include estimating, by the machine learning model 65052 relationships between a dependent variable, i.e., an outcome variable, and one or more independent variables, i.e., predictors, covariates, and/or features. Similarity learning may include learning, by the machine learning model 65052, from examples using a similarity function, the similarity function being designed to measure how similar or related two objects are.

[2512] In example embodiments, the machine learning model 65052 may be defined via unsupervised learning, i.e., one or more algorithms configured to build a mathematical model of a set of data containing only inputs by finding structure in the data such as grouping or clustering of data points. In example embodiments, the machine learning model 65052 may LEARN from test data, i.e., training data, that has not been labeled, classified, or categorized. The unsupervised learning algorithm may include identifying, by the machine learning model 65052, commonalities in the training data and learning by reacting based on the presence or absence of the identified commonalities in new pieces of data. In example embodiments, the machine learning model 65052 may generate one or more probability density functions. In example embodiments, the machine learning model 65052 may LEARN by performing cluster analysis, such as by assigning a set of observations into subsets, i.e., clusters, according to one or more predesignated criteria, such as according to a similarity metric of which internal compactness, separation, estimated density, and/or graph connectivity are factors.

[2513] In example embodiments, the machine learning model 65052 may be defined via semi-supervised learning, i.e., one or more algorithms using training data wherein some training examples may be missing training labels. The semi-supervised learning may be weakly supervised learning, wherein the training labels may be noisy, limited, and/or imprecise. The noisy, limited, and/or imprecise training labels may be cheaper and/or less labor intensive to produce, thus allowing the machine learning model 65052 to train on a larger set of training data for less cost and/or labor.

[2514] In example embodiments, the machine learning model 65052 may be defined via reinforcement learning, such as one or more algorithms using dynamic programming techniques such that the machine learning model 65052 may train by taking actions in an environment in order to maximize a cumulative reward. In example embodiments, the training data is represented as a Markov Decision Process.

[2515] In example embodiments, the machine learning model 65052 may be defined via self-learning, wherein the machine learning model 65052 is configured to train using training data with no external rewards and no external teaching, such as by employing a Crossbar Adaptive Array (CAA). The CAA may compute decisions about actions and/or emotions about consequence situations in a crossbar fashion, thereby driving teaching of the machine learning model 65052 by interactions between cognition and emotion.

[2516] In example embodiments, the machine learning model 65052 may be defined via feature learning, i.e., one or more algorithms designed to discover increasingly accurate and/or apt representations of one or more inputs provided during training, e.g., training data. Feature learning

may include training via principal component analysis and/or cluster analysis. Feature learning algorithms may include attempting, by the machine learning model 65052, to preserve input training data while also transforming the input training data such that the transformed input training data is useful. In example embodiments, the machine learning model 65052 may be configured to transform the input training data prior to performing one or more classifications and/or predictions of the input training data. Thus, the machine learning model 65052 may be configured to reconstruct input training data from one or more unknown data-generating distributions without necessarily conforming to implausible configurations of the input training data according to the distributions. In example embodiments, the feature learning algorithm may be performed by the machine learning model 65052 in a supervised, unsupervised, or semi-supervised manner.

[2517] In example embodiments, the machine learning model 65052 may be defined via anomaly detection, i.e., by identifying rare and/or outlier instances of one or more items, events and/or observations. The rare and/or outlier instances may be identified by the instances differing significantly from patterns and/or properties of a majority of the training data. Unsupervised anomaly detection may include detecting of anomalies, by the machine learning model 65052, in an unlabeled training data set under an assumption that a majority of the training data is "normal." Supervised anomaly detection may include training on a data set wherein at least a portion of the training data has been labeled as "normal" and/or "abnormal."

[2518] In example embodiments, the machine learning model 65052 may be defined via robot learning. Robot learning may include generation, by the machine learning model 65052, of one or more curricula, the curricula being sequences of learning experiences, and cumulatively acquiring new skills via exploration guided by the machine learning model 65052 and social interaction with humans by the machine learning model 65052. Acquisition of new skills may be facilitated by one or more guidance mechanisms such as active learning, maturation, motor synergies, and/or imitation.

[2519] In example embodiments, the machine learning model 65052 can be defined via association rule learning. Association rule learning may include discovering relationships, by the machine learning model 65052, between variables in databases, in order to identify strong rules using some measure of "interestingness." Association rule learning may include identifying, learning, and/or evolving rules to store, manipulate and/or apply knowledge. The machine learning model 65052 may be configured to learn by identifying and/or utilizing a set of relational rules, the relational rules collectively representing knowledge captured by the machine learning model 65052. Association rule learning may include one or more of learning classifier systems, inductive logic programming, and artificial immune systems. Learning classifier systems are algorithms that may combine a discovery component, such as one or more genetic algorithms, with a learning component, such as one or more algorithms for supervised learning, reinforcement learning, or unsupervised learning. Inductive logic programming may include rule-learning, by the machine learning model 65052, using logic programming to represent one or more of input examples, background knowledge, and hypothesis determined by the machine learning model 65052 during training. The machine learning model 65052 may be configured to derive a hypothesized logic

program entailing all positive examples given an encoding of known background knowledge and a set of examples represented as a logical database of facts.

[2520] In example embodiments, the platform can deploy many systems and methods for the industrial internet of things (IIoT) including solutions that can be configured as IIoT in a Box and other system configurations for IIoT; IIoT interface devices and systems (e.g., AR, VR, xR, wearables, and the like); advanced chips, boards, and switches for IIoT applications and the like. In example embodiments, the platform can deploy many different systems and methods for data collection, sensor fusion, data management and artificial intelligence; systems and methods for intelligent data collection for IIoT; systems and methods for equipment-specific data collection and management systems; systems and methods for biology-based data management for IIoT; systems and methods for advanced visual/optical sensing for IIoT intelligence; systems and methods for sensor fusion and sensor package configuration for IIoT intelligence; systems and methods for smart data pipelines for IIoT storage and computation; systems and methods for advanced, coordinated data collection and operations systems (e.g., drones, robotics, and the like); and systems and methods for advanced vibration sensing, monitoring and diagnostics. In example embodiments, the platform can deploy many systems and methods for advanced operational awareness and control including systems and methods for advanced industrial process control (e.g., hydrolysis to produce hydrogen for industrial heating, cooking, processing, etc.); systems and methods for artificial intelligence and data processing for detection and prediction of IIoT patterns and states; systems including platforms and associated methodologies for agile management and governance of IIoT operations (e.g., twins; dashboards; policy engine and the like); systems and methods for domain-specific applications of IIoT intelligence platform (e.g., oil & gas; mining, agricultural, municipal, and the like); systems and methods for converged IIoT platforms; and systems and methods for automated industrial service ecosystems. In example embodiments, the platform can deploy many networking and computation for IIoT entities including systems and methods for convergence of edge and networking; systems and methods for enhancement of radio frequency (RF) networking for IIoT; systems and methods for quantum algorithms in combination with artificial intelligence for IIoT intelligence; and systems and methods for smart networking protocols.

[2521] FIG. 254 illustrates an example quantum computing system 67000 according to some embodiments of the present disclosure. In example embodiments, the quantum computing system 67000 provides a framework for providing a set of quantum computing services to one or more quantum computing clients. In example embodiments, the quantum computing system 67000 framework may be at least partially replicated in respective quantum computing clients (e.g., various industrial entities). In these embodiments, an individual client may include some or all of the capabilities of the quantum computing system 67000, whereby the quantum computing system 67000 is adapted for the specific functions performed by the subsystems of the quantum computing client. Additionally, or alternatively, in example embodiments, the quantum computing system 67000 may be implemented as a set of microservices, such that different quantum computing clients may leverage the quantum computing system 67000 via one or more APIs exposed to the

quantum computing clients. In these embodiments, the quantum computing system 67000 may be configured to perform various types of quantum computing services that may be adapted for different quantum computing clients. In either of these configurations, a quantum computing client may provide a request to the quantum computing system 67000, whereby the request is to perform a specific task (e.g., an optimization). In response, the quantum computing system 67000 executes the requested task and returns a response to the quantum computing client.

[2522] Referring to FIG. 254, in example embodiments, the quantum computing system 67000 may include a quantum adapted services library 67002, a quantum general services library 67004, a quantum data services library 67006, a quantum computing engine library 67008, a quantum computing configuration service 67010, a quantum computing execution system 67012, and quantum computing API interface 67014.

[2523] In example embodiments, the quantum computing engine library 67008 includes quantum computing engine configurations 67016 and quantum computing process modules 67018 based on various supported quantum models. In example embodiments, the quantum computing system 67000 may support many different quantum models, including, but not limited to, the quantum circuit model, quantum Turing machine, spintronic computing system (such as using spin-orbit coupling to generate spin-polarized electronic states in non-magnetic solids, such as ones using diamond materials), adiabatic quantum computer, one-way quantum computer, quantum annealing, and various quantum cellular automata. Under the quantum circuit model, quantum circuits may be based on the quantum bit, or "qubit", which is somewhat analogous to the bit in classical computation. Qubits may be in a 1 or 0 quantum state or they may be in a superposition of the 1 and 0 states. However, when qubits have measured the result of a measurement, qubits will always be in is always either a 1 or 0 quantum state. The probabilities related to these two outcomes depend on the quantum state that the qubits were in immediately before the measurement. Computation is performed by manipulating qubits with quantum logic gates, which are somewhat analogous to classical logic gates.

[2524] In example embodiments, the quantum computing system 67000 may be physically implemented using an analog approach or a digital approach. Analog approaches may include, but are not limited to, quantum simulation, quantum annealing, and adiabatic quantum computation. In example embodiments, digital quantum computers use quantum logic gates for computation. Both analog and digital approaches may use quantum bits, or qubits.

[2525] In example embodiments, the quantum computing system 67000 includes a quantum annealing module 67020 wherein the quantum annealing module may be configured to find the global minimum or maximum of a given objective function over a given set of candidate solutions (e.g., candidate states) using quantum fluctuations. As used herein, quantum annealing may refer to a meta-procedure for finding a procedure that identifies an absolute minimum or maximum, such as a size, length, cost, time, distance or other measure, from within a possibly very large, but finite, set of possible solutions using quantum fluctuation-based computation instead of classical computation. The quantum annealing module 67020 may be leveraged for problems where the

search space is discrete (e.g., combinatorial optimization problems) with many local minima, such as finding the ground state of a spin glass or the traveling salesman problem.

[2526] In example embodiments, the quantum annealing module 67020 starts from a quantum-mechanical superposition of all possible states (candidate states) with equal weights. The quantum annealing module 67020 may then evolve, such as following the time-dependent Schrödinger equation, a natural quantum-mechanical evolution of systems (e.g., physical systems, logical systems, or the like). In example embodiments, the amplitudes of all candidate states change, realizing quantum parallelism according to the time-dependent strength of the transverse field, which causes quantum tunneling between states. If the rate of change of the transverse field is slow enough, the quantum annealing module 67020 may stay close to the ground state of the instantaneous Hamiltonian. If the rate of change of the transverse field is accelerated, the quantum annealing module 67020 may leave the ground state temporarily but produce a higher likelihood of concluding in the ground state of the final problem energy state or Hamiltonian.

[2527] In example embodiments, the quantum computing system 67000 may include arbitrarily large numbers of qubits and may transport ions to spatially distinct locations in an array of ion traps, building large, entangled states via photonically connected networks of remotely entangled ion chains.

[2528] In some implementations, the quantum computing system 67000 includes a trapped ion computer module 67022, which may be a quantum computer that applies trapped ions to solve complex problems. Trapped ion computer module 67022 may have low quantum decoherence and may be able to construct large solution states. Ions, or charged atomic particles, may be confined and suspended in free space using electromagnetic fields. Qubits are stored in stable electronic states of each ion, and quantum information may be transferred through the collective quantized motion of the ions in a shared trap (interacting through the Coulomb force). Lasers may be applied to induce coupling between the qubit states (for single-qubit operations) or coupling between the internal qubit states and the external motional states (for entanglement between qubits).

[2529] In some embodiments of the invention, a traditional computer, including a processor, memory, and a graphical user interface (GUI), may be used for designing, compiling, and providing output from the execution and the quantum computing system 67000 may be used for executing the machine language instructions. In some embodiments of the invention, the quantum computing system 67000 may be simulated by a computer program executed by the traditional computer. In such embodiments, a superposition of states of the quantum computing system 67000 can be prepared based on input from the initial conditions. Since the initialization operation available in a quantum computer can only initialize a qubit to either the $|0\rangle$ or $|1\rangle$ state, initialization to a superposition of states is physically unrealistic. For simulation purposes, however, it may be useful to bypass the initialization process and initialize the quantum computing service 67000 directly.

[2530] In example embodiments, the quantum computing system 67000 provides various quantum data services, including quantum input filtering, quantum output filtering, quantum application filtering, and a quantum database engine.

[2531] In example embodiments, the quantum computing system 67000 may include a quantum input filtering service 67024. In example embodiments, quantum input filtering service 67024 may be configured to select whether to run a model on the quantum computing system 67000 or to run the model on a classic computing system. In example embodiments, quantum input filtering service 67024 may filter data for later modeling on a classic computer. In example embodiments, the quantum computing system 67000 may provide input to traditional compute platforms while filtering out unnecessary information from flowing into distributed systems. In example embodiments, the platform 67000 may trust through filtered specified experiences for intelligent agents.

[2532] In example embodiments, an industrial internet of things system may include model or system for automatically determining, based on a set of inputs, whether to deploy quantum computational or quantum algorithmic resources to an industrial activity, whether to deploy traditional computational resources and algorithms, or whether to apply a hybrid or combination of them. In example embodiments, inputs to a model or automation system may include demand information, supply information, energy cost information, capital costs for computational resources, development costs (such as for algorithms), energy costs, operational costs (including labor and other costs), performance information on available resources (quantum and traditional), and any of the many other data sets that may be used to simulate (such as using any of a wide variety of simulation techniques described herein and/or in the documents incorporated herein by refence) and/or predict the difference in outcome between a quantum-optimized result and a non-quantum-optimized result. A machine learned model (including in a DPANN system) may be trained, such as by deep learning on outcomes or by a data set from human expert decisions, to determine what set of resources to deploy given the input data for a given request. The model may itself be deployed on quantum computational resources and/or may use quantum algorithms, such as quantum annealing, to determine whether, where and when to use quantum systems, conventional systems, and/or hybrids or combinations.

[2533] In some embodiments of the invention, the quantum computing system 67000 may include a quantum output filtering service 67026. In example embodiments, the quantum output filtering service 67026 may be configured to select a solution from solutions of multiple neural networks. For example, multiple neural networks may be configured to generate solutions to a specific problem and the quantum output filtering service 67026 may select the best solution from the set of solutions.

[2534] In example embodiments, the quantum computing system 67000 connects and directs a neural network development or selection process. In this embodiment, the quantum computing system 67000 may directly program the weights of a neural network such that the neural network gives the desired outputs. This quantum-programmed neural network may then operate without the oversight of the quantum computing system 67000 but will still be operating within the expected parameters of the desired computational engine.

[2535] In example embodiments, the quantum computing system 67000 includes a quantum database engine 67028. In example embodiments, the quantum database engine 67028 is

configured with in-database quantum algorithm execution. In example embodiments, a quantum query language may be employed to query the quantum database engine 67028. In example embodiments, the quantum database engine may have an embedded policy engine 67030 for prioritization and/or allocation of quantum workflows, including prioritization of query workloads, such as based on overall priority as well as the comparative advantage of using quantum computing resources versus others. In example embodiments, quantum database engine 67028 may assist with the recognition of entities by establishing a single identity for that is valid across interactions and touchpoints. The quantum database engine 67028 may be configured to perform optimization of data matching and intelligent traditional compute optimization to match individual data elements. The quantum computing system 67000 may include a quantum data obfuscation system for obfuscating data.

[2536] The quantum computing system 67000 may include, but is not limited to, analog quantum computers, digital computers, and/or error-corrected quantum computers. Analog quantum computers may directly manipulate the interactions between qubits without breaking these actions into primitive gate operations. In example embodiments, quantum computers that may run analog machines include, but are not limited to, quantum annealers, adiabatic quantum computers, and direct quantum simulators. The digital computers may operate by carrying out an algorithm of interest using primitive gate operations on physical qubits. Error-corrected quantum computers may refer to a version of gate-based quantum computers made more robust through the deployment of quantum error correction (QEC), which enables noisy physical qubits to emulate stable logical qubits so that the computer behaves reliably for any computation. Further, quantum information products may include, but are not limited to, computing power, quantum predictions, and quantum inventions.

[2537] In example embodiments, the quantum computing system 67000 is configured as an engine that may be used to optimize traditional computers, integrate data from multiple sources into a decision-making process, and the like. The data integration process may involve real-time capture and management of interaction data by a wide range of tracking capabilities, both directly and indirectly related to industrial activities. In example embodiments, the quantum computing system 67000 may be configured to accept cookies, email addresses and other contact data, social media feeds, news feeds, event and transaction log data (including transaction events, network events, computational events, and many others), event streams, results of web crawling, distributed ledger information (including blockchain updates and state information), results from distributed or federated queries of data sources, streams of data from chat rooms and discussion forums, and many others.

[2538] In example embodiments, the quantum computing system 67000 includes a quantum register having a plurality of qubits. Further, the quantum computing system 67000 may include a quantum control system for implementing the fundamental operations on each of the qubits in the quantum register and a control processor for coordinating the operations required.

[2539] In example embodiments, the quantum computing system may include a quantum-enabled or other risk identification module that is configured to perform risk identification and/or

mitigation. The steps that may be taken by the risk identification module may include, but are not limited to, risk identification, impact assessment, and the like. In example embodiments, the risk identification module determines a risk type from a set of risk types. In example embodiments, risks may include, but are not limited to, preventable, strategic, and external risks. Preventable risks may refer to risks that come from within and that can usually be managed on a rule-based level, such as employing operational procedures monitoring and employee and manager guidance and instruction. Strategy risks may refer to those risks that are taken on voluntarily to achieve greater rewards. External risks may refer to those risks that originate outside and are not in the businesses' control (such as natural disasters). External risks are not preventable or desirable. In example embodiments, the risk identification module can determine a predicted cost for any category of risk. The risk identification module may perform a calculation of current and potential impact on an overall risk profile. In example embodiments, the risk identification module may determine the probability and significance of certain events. Additionally, or alternatively, the risk identification module may be configured to anticipate events.

[2540] In example embodiments, the quantum computing system 67000 or other system of the platform 67000 is configured for accelerated sampling from stochastic processes for risk analysis. In example embodiments, quantum-simulated accelerated testing is initialized to hold accelerated life tests with constant-stress loadings, including accelerated degradation tests and time-varying stress loadings.

[2541] In example embodiments, the quantum computing system 67000 includes a quantum prediction module, which is configured to generate predictions. Furthermore, the quantum prediction module may construct classical prediction engines to further generate predictions, reducing the need for ongoing quantum calculation costs, which, can be substantial compared to traditional computers.

[2542] In example embodiments, the quantum computing system 67000 may include a quantum principal component analysis (QPCA) algorithm that may process input vector data if the covariance matrix of the data is efficiently obtainable as a density matrix, under specific assumptions about the vectors given in the quantum mechanical form. It may be assumed that the user has quantum access to the training vector data in a quantum memory. Further, it may be assumed that each training vector is stored in the quantum memory in terms of its difference from the class means. These QPCA can then be applied to provide for dimension reduction using the calculational benefits of a quantum method.

[2543] In example embodiments, the quantum computing system 67000 includes a quantum continual learning (QCL) system 67032, wherein the QCL system 67032 learns continuously and adaptively about the external world, enabling the autonomous incremental development of complex skills and knowledge by updating a quantum model to account for different tasks and data distributions. The QCL system 67032 operates on a realistic time scale where data and/or tasks become available only during operation. Previous quantum states can be superimposed into the quantum engine to provide the capacity for QCL. Because the QCL system 67032 is not constrained to a finite number of variables that can be processed deterministically, it can

continuously adapt to future states, producing a dynamic continual learning capability. The QCL system 67032 may have applications where data distributions stay relatively static, but where data is continuously being received. For example, the QCL system 67032 may be used in quantum recommendation applications or quantum anomaly detection systems where data is continuously being received and where the quantum model is continuously refined to provide for various outcomes, predictions, and the like. QCL enables asynchronous alternate training of tasks and only updates the quantum model on the real-time data available from one or more streaming sources at a particular moment.

[2544] In example embodiments, the QCL system 67032 operates in a complex environment in which the target data keeps changing based on a hidden variable that is not controlled. In example embodiments, the QCL system 67032 can scale in terms of intelligence while processing increasing amounts of data and while maintaining a realistic number of quantum states. The QCL system 67032 applies quantum methods to drastically reduce the requirement for storage of historic data while allowing the execution of continuous computations to provide for detail-driven optimal results. In example embodiments, a QCL system 67032 is configured for unsupervised streaming perception data since it continually updates the quantum model with new available data.

[2545] In example embodiments, QCL system 67032 enables multi-modal-multi-task quantum learning. The QCL system 67032 is not constrained to a single stream of perception data but allows for many streams of perception data from different sensors and input modalities. In example embodiments, the QCL system 67032 can solve multiple tasks by duplicating the quantum state and executing computations on the duplicate quantum environment. A key advantage to QCL is that the quantum model does not need to be retrained on historic data, as the superposition state holds information relating to all prior inputs. Multi-modal and multi-task quantum learning enhance quantum optimization since it endows quantum machines with reasoning skills through the application of vast amounts of state information.

[2546] In example embodiments, the quantum computing system 67000 supports quantum superposition, or the ability of a set of states to be overlaid into a single quantum environment.

[2547] In example embodiments, the quantum computing system 67000 supports quantum teleportation. For example, information may be passed between photons on chipsets even if the photons are not physically linked.

[2548] In example embodiments, quantum-aware device stacks may be supported, including quantum-aware device-level kits, quantum-aware industrial Internet of Things (IoT) kits, quantum-enabled FPGAs, and systems with awareness of capabilities of different quantum computer types and/or different quantum algorithm types.

[2549] In example embodiments, the quantum computing system 67000 may leverage one or artificial networks to fulfill the request of a quantum computing client. For example, the quantum computing system 67000 may leverage a set of artificial neural networks to identify patterns in images (e.g., using image data from a liquid lens system), perform binary matrix factorization, perform topical content targeting, perform similarity-based clustering, perform collaborative filtering, perform opportunity mining, or the like.

[2550] In example embodiments, an information technology system may include a hybrid computing allocation system for prioritization and allocation of quantum computing resources and traditional computing resources. In example embodiments, the prioritization and allocation of quantum computing resources and traditional computing resources may be measure-based (e.g., measuring the extent of the advantage of the quantum resource relative to other available resources), cost-based, optimality-based, speed-based, impact-based, or the like. In some embodiments the hybrid computing allocation system is configured to perform time-division multiplexing between the quantum computing system 67000 and a traditional computing system. In example embodiments, the hybrid computing allocation system may automatically track and report on the allocation of computational resources, the availability of computational resources, the cost of computational resources, and the like.

[2551] In example embodiments, the quantum computing system 67000 may be leveraged for queue optimization for utilization of quantum computing resources, including context-based queue optimizations.

[2552] In example embodiments, the quantum computing system 67000 may support quantum-computation-aware location-based data caching.

[2553] In example embodiments, the quantum computing system 67000 may be leveraged for optimization of various system resources, including the optimization of quantum computing resources, traditional computing resources, energy resources, human resources, robotic fleet resources, smart container fleet resources, I/O bandwidth, storage resources, network bandwidth, attention resources, or the like.

[2554] The quantum computing system 67000 may be implemented in a manner in which a complete range of capabilities are available to or as part of any configured service. Configured quantum computing services may be configured with subsets of these capabilities to perform specific predefined functions, produce newly defined functions, or various combinations of both.

[2555] FIG. 255 illustrates quantum computing service request handling according to some embodiments of the present disclosure. A directed quantum computing request 67102 may come from one or more quantum-aware devices or stack of devices, where the request is for known application configured with specific quantum instance(s), quantum computing engine(s), or other quantum computing resources, and where data associated with the request may be preprocessed or otherwise optimized for use with quantum computing.

[2556] A general quantum computing request 67104 may come from any system or configured service, where the requestor has determined that quantum computing resources may provide additional value or other improved outcomes. Improved outcomes may also be suggested by the quantum computing service in association with some form of monitoring and analysis. For a general quantum computing request 67104, input data may not be structured or formatted as necessary for quantum computing.

[2557] In example embodiments, external data requests 67106 may include any available data that may be necessary for training new quantum instances. The sources of such requests could be public data, sensors, ERP systems, and many others.

[2558] Incoming operating requests and associated data may be analyzed using a standardized approach that identifies one or more possible sets of known quantum instances, quantum computing engines, or other quantum computing resources that may be applied to perform the requested operation(s). Potential existing sets may be identified in the quantum set library 67108.

[2559] In example embodiments, the quantum computing system 67000 includes a quantum computing configuration service 67010. The quantum computing configuration service may work alone or with the intelligence service 67034 to select a best available configuration using a resource and priority analysis that also includes the priority of the requestor. The quantum computing configuration service may provide a solution (YES) or determine that a new configuration is required (NO).

[2560] In one example, the requested set of quantum computing services may not exist in the quantum set library 67108. In this example, one or more new quantum instances must be developed (trained) using available data. For example, a quantum computing module for optimizing the performance of a specific machine may exist in the quantum set library 67108. However, requestor inputs identified the need to optimize a different model of the machine. In this case, quantum instance training may work with the intelligence service 67034 to train a new instance for the different machine model using a range of public data such as machine reviews, technical specifications, and so forth. In example embodiments, alternate configurations may be developed with assistance from the intelligence service 67034 to identify alternate ways to provide all or some of the requested quantum computing services until appropriate resources become available. For example, a quantum/traditional hybrid model may be possible that provides the requested service, but at a slower rate.

[2561] In example embodiments, alternate configurations may be developed with assistance from the intelligence service 67034 to identify alternate and possibly temporary ways to provide all or some of the requested quantum computing services. For example, a hybrid quantum/traditional model may be possible that provides the requested service, but at a slower rate. This may also include a feedback learning loop to adjust services in real time or to improve stored library elements.

[2562] When a quantum computing configuration has been identified and available, it is allocated and programmed for execution and delivery of one or more quantum states (solutions).

Biological Systems for Industrial Internet of Things

[2563] Techniques described herein improve the ability of networks and systems to collect, transmit, and process large volumes of data, especially data from sensors and other industrial internet of things data generators. These techniques include using a thalamus service that provides an equivalent to a biological thalamus, a neural system for filtering and relaying data. The thalamus service described herein can receive large volumes of information and quickly prioritize the information, passing on the most importing information so that limited transmission, processing, collection, and/or analysis resources are not overwhelmed by volume of incoming information.

[2564] Additionally, a predictive model communication protocol (PMCP) is described herein. PMCP may be used to reduce a volume of transmitted data, especially when the data is predictable

or usually predictable. PMCP may operate by transmitting predictive model parameters instead of some or all of the data values that would normally be transmitted by a sensor device or other data source. For example, a device implementing PMCP may continually receive inputs (e.g., sensor data) and train a predictive model using the stream of sensor data. Rather than transmitting the sensor data, which may use significant network and/or processing resources, the PMCP device may transmit the model parameters, which may be used by a receiving device to operate a predictive model to predict current and future sensor data. Thus, the receiving device may have a predictive model of sensor data without receiving the sensor data. In example embodiments, if the sensor data at the PMCP device begins operating outside of expectations, the model parameters may be re-transmitted to the receiving device, which may update its predictive model and thereby obtain more accurate predictive data.

[2565] In example embodiments, to optimize decision-making, quantum computers and/or predictive models may be used with the techniques described herein. Furthermore, quantum coordination can be applied to allow for disparate units to securely coordinate actions (e.g., without the need for traditional communication mechanisms). Accordingly, techniques described herein may use a combination of decentralized biology-based decision-making capabilities distributed throughout devices within the industrial internet of things and quantum capabilities.

[2566] FIG. 256 shows a thalamus service 68000 and a set of input sensors streaming data from various sources across the industrial internet of things and the control system 68002 with its centrally-managed data sources 68004. The thalamus service 68000 filters the inputs from the various data sources 68004 into the control system 68002 such that the control system is never overwhelmed by the total volume of information. In example embodiments, the thalamus service 68000 provides an information suppression mechanism for information flows. This mechanism monitors all data streams and suppresses and/or filters irrelevant data streams by ensuring that the maximum data flows from all input sensors are always constrained.

[2567] In example embodiments, the thalamus service 68000 may be a gateway for all communication that responds to the prioritization of the control system 68002. The control system 68002 may decide to change the prioritization of the data streamed from the thalamus service 68000, for example, during a known fire in an isolated area, and the event may direct the thalamus service 68000 to continue to provide flame sensor information despite the fact that majority of this data is not unusual. The thalamus service 68000 may be an integral part of the overall communication framework.

[2568] In example embodiments, the thalamus service 68000 includes an intake management system 68006. The intake management system 68006 may be configured to receive and process multiple large datasets by converting them into data streams that are sized and organized for subsequent use by a central control system 68002. For example, a robot may include vision and sensing systems that are used by the central control system 68002 (which may be on-board the robot and/or in a separate device in communication with the robot) to identify and move through an environment in real-time. The intake management system 68006 can facilitate robot decision-making by parsing, filtering, classifying, or otherwise reducing the size and increasing the utility

of multiple large datasets that would otherwise overwhelm the central control system 68002. In example embodiments, the intake management system may include an intake controller 68008 that works with the intelligence service 68010 to evaluate incoming data and take actions-based evaluation results. Evaluations and actions may include specific instruction sets received by the thalamus service 68000, for example the use of a set of specific compression and prioritization tools stipulated within a "Networking" library module. In another example, thalamus service inputs may direct the use of specific filtering and suppression techniques. In a third example, thalamus service inputs may stipulate data filtering associated with an area of interest such as a certain type of machine. The intake management system is also configured to recognize and manage datasets that are in a vectorized format such as in accordance with a predictive model communication protocol (PMCP) (discussed below), where the datasets may be passed directly to the central control system 68002, or alternatively deconstructed and processed separately. The intake management system 68006 may include a learning module that receives data from external sources that enables improvement and creation of application and data management library modules. In some cases, the intake management system 68006 may request external data to augment existing datasets.

[2569] In example embodiments, the control system 68002 may direct the thalamus service 68000 to alter its filtering to provide more input from a set of specific sources. This indication to provide more input is handled by the thalamus service 68000. For example, the thalamus service may suppress other information flows to constrain the total data flows to within a volume that the central control system can handle.

[2570] In example embodiments, the thalamus service 68000 can operate by suppressing data based on several different factors including zero or more default factors. For example, in some embodiments the default factors may include an "unusualness factor" that may be a value that indicates a divergence or a degree of divergence of the data from an expected dataset. In example embodiments, the unusualness factor is constantly monitored for all inputs or some of the inputs (e.g., some of the input sensors).

[2571] In example embodiments, the thalamus service 68000 may suppress data based on geospatial factors. Examples of geospatial factors may include location data, motion data, acceleration data, vibration data, and/or any other data indicating an absolute or relative location, change in location over time, other derivatives or integrals of location over time, etc. The thalamus service 68000 may be aware of the geospatial factors for some or all of the sensors and thus is able to look for unusual patterns in data based on geospatial context and suppress data accordingly.

[2572] In example embodiments, the thalamus service 68000 may suppress data based on temporal factors. Data can be suppressed temporally, for example, if the cadence of the data can be reduced such that the overall data stream is filtered to a level that can be handled by the control system 68002 and/or a central processing unit.

[2573] In example embodiments, the thalamus service 68000 may suppress data based on contextual factors. In example embodiments, context-based filtering is a filtering event in which the thalamus service 68000 is aware of some context-based event. Context-based events, for

example, may include one or more notifications of unusual behavior by other sensors or systems (which may lead to temporary suppression of less important data), one or more human inputs (e.g., a human disabling a security alert, which may suppress a previous focus on security data), one or more events triggered by other systems or sensors (e.g., an automated security alert, which may lead to suppression of certain data to allow resources to be dedicated to security data collection, transmission, and analysis), one or more contexts detected from other sensor data (e.g., a reduction in available bandwidth reported by a network sensor, which may lead to the suppression of certain data until available bandwidth improves), or any other context-based condition or event. In this context, the filtering may suppress information flows not relating to the data from the event.

[2574] In example embodiments, the thalamus service 68000 may receive data from a variety of data sources 68004, including analyses 68018, databases 68020, sensors 68022, and/or reports 68024. For example, the thalamus service 68000 may receive analyses and/or reports from other analysis/processing/reporting devices that have already pre-processed sensor data or other data. Additionally or alternatively, the thalamus service 68000 may receive data (e.g., historical data) that is stored in a database 68020 in addition to current or historical data from sensors 68022. In example embodiments, data may be received and/or generated (e.g., predictive models may generate future data) from the PMCP device interface 68052.

[2575] In example embodiments, the thalamus service 68000 may process and/or interpret inputs from any of the data sources 68004 based on an intake application library 68012, which may include a networking library 68014, a security library 68016, and/or any other library for interpreting various types of input data. For example, the thalamus service 68000 may use a networking 68014 library to parse, interpret, extract, and/or otherwise process network data (e.g., data received from networking sensors or devices, networking analyses, networking reports, network database data, etc.). Similarly, the thalamus service 68000 may use a security 68016 library to parse, interpret, extract, and/or otherwise process security data (e.g., data received from security sensors or devices, security analyses, security reports, security database data, etc.). In example embodiments, the intake data may also be processed using an intake learning module 68026, which may use one or more artificial intelligence techniques to pre-process the data, generate predictive models using the data, predict future states of the data, and/or the like. After processing using the intake application library 68012 and/or the intake learning module 68026, the data may be ready for management by the intake data management system 68028.

[2576] The intake data management system 68028 may process the data by prioritizing 68030, formatting 68032, suppressing 68034, using an area focus 68036, filtering 68038, and/or combining 68040 the data. The prioritizing 68030 may involve ranking or otherwise assigning priorities (e.g., categories, numerical priority scores, etc.) such that limited resources may be assigned to the most important data. For example, the suppressing 68034 and/or filtering 68038 may operate based on priorities in order to suppress or filter out the least important data (e.g., the data associated with a lowest priority score) in order to avoid overwhelming limited transmission, processing, and/or analysis resources. The formatting 68032 may involve formatting data in order to allow for easier management, which may involve compressing or otherwise dropping certain

parts of data to reduce use of transmission resources, un-compressing data to reduce use of decompression resources (e.g., if bandwidth is sufficient and data is important), formatting data to emphasize or de-emphasize certain aspects, or otherwise adjust formatting. In example embodiments, the formatting 68032 may depend on the prioritizing 68030 such that more important data may be formatted in order to allow for more or better analysis, while less important data may be formatted in order to reduce it usage of various resources.

[2577] In example embodiments, the suppressing 68034 may involve reducing the amount of data, the number of destinations to which the data is transmitted, or otherwise reducing the usage of limited resources (e.g., bandwidth, processing, analysis, etc.) of the data. In example embodiments, suppressed data may be stored (e.g., in a database) and dealt with (e.g., transmitted, processed) at a later time. In example embodiments, the suppressing 68034 may be based various factors as described above.

[2578] In example embodiments, an area focus 68036 may involve increasing the attention paid to certain high priority data. For example, during a security incident, security sensor data may be sent to additional destinations, processed using additional analyses, allowed additional bandwidth and processing power, and/or the like. In example embodiments, an area focus 68036 may cause the suppression or filtering of other data that is not associated with the area focus 68036.

[2579] In example embodiments, the filtering 68038 may involve ignoring, deleting, or otherwise removing data that is not important (e.g., does not match an area focus 68036, is low priority, etc.). In example embodiments, data may be initially suppressed (e.g., reduced or stored for later), but conditions may further change, causing the data to be filtered (e.g., deleted, ignored). Thus, intake data management system 68028 may allow for a progressive downgrade of data by first suppressing and later filtering the data depending on conditions.

[2580] In example embodiments, the combining 68040 may include combining various types of data in order to provide better analyses, generate new data, reduce the volume of data (e.g., by combining multiple data values into a single data value), improve the quality of data (e.g., by averaging different sensor readings to obtain a more accurate average reading), and/or the like. In example embodiments, lower priority data may be combined with other data in order to reduce requirements. Additionally or alternatively, higher priority data may be combined with other data in order to improve data quality.

[2581] In example embodiments, the intake data management system may interface with an intake controller 68008 and/or an intelligence system 68042. The intelligence system 68042, for example, may use various artificial intelligence techniques to perform the intake data management (e.g., prioritize the data, format the data, suppress the data, select an area focus and/or assign data to an area focus, filter the data, combine the data, etc.), predict the outcomes of intake data management, predict future data values, and/or the like. Additionally or alternatively, the intake controller 68008 and/or an intelligence system 68042 may operate in accordance with configured thalamus parameters 68044, which may govern the intake data management, the artificial intelligence techniques (e.g., the parameters may be model parameters for AI models), and/or otherwise configure the operations of the intake management system 68006.

[2582] In example embodiments, the control system 68002 may, in some cases, use a quantum computing service 68046, which may provide quantum computing resources to more quickly process large volumes of data, use quantum models, and/or the like.

[2583] The control system 68002 may further comprise one or more data interfaces 68048 for receiving data from various data sources 68004 and transmitting the data (e.g., after intake data management) to various destinations. In example embodiments, the control system 68002 may include other system subsystems 68050, such as analysis subsystems, various processing chips, or any other subsystems that may use the managed data to make decisions, generate analyses, or otherwise perform data operations. In example embodiments, an intelligence service 68010 may operate to route the managed data to various other system subsystems 68050, or otherwise perform initial and/or final processing on the data.

[2584] In example embodiments, the control system 68002 can override the thalamus filtering and decide to focus on a different area for any specific reason. For example, during a security incident, the control system may route around thalamus filtering (which might normally de-prioritize data from security sensors) in order to ensure that data from security sensors is delivered in full without any de-prioritization, suppression, filtering, etc. As another example, during regular inspections of equipment, sensor data that measures operation of the equipment (e.g., vibration sensor data) may be un-suppressed, even if the data appears to be within normal parameters and therefore might usually be suppressed or filtered.

[2585] In example embodiments, the control system 68002 may include a PMCP device interface 68052, which may be used to transmit and/or receive data using PMCP. Details of a PMCP device interface are further shown within a second PMCP device interface 68060. In example embodiments, the PMCP device interface 68052 may be in communication with the PMCP device interface 68060. The PMCP device interface 68052 may have the same components as shown within the PMCP device interface 68060.

[2586] In example embodiments, the PMCP device interface may be used to convert data to a vectorized format prior to transmission. In these embodiments, a vector may be considered an example of a simple predictive model (e.g., a vector may indicate an amount of change and a direction of change for a data value, thus predicting a future state of a data value if the change continues). For example, the conversion of a long sequence of oftentimes similar data values into a vector indicating an amount and direction of change, which may imply a future state of the data values, makes the communication of the data values both smaller in size and forward looking in nature.

[2587] In example embodiments, PMCP may use various types of predictive models to predict current and future data values, including weighted moving average; Kalman filtering; exponential smoothing; autoregressive moving average (ARMA) (forecasts depend on past values of the variable being forecast, and on past prediction errors); autoregressive integrated moving average (ARIMA) (ARMA on the period-to-period change in the forecasted variable); extrapolation; linear prediction; trend estimation (predicting the variable as a linear or polynomial function of time); growth curve (e.g., statistics); and recurrent neural network based forecasting.

[2588] Using the PMCP protocol, instead of traditional streams where individual data items are transmitted, vectors representing how the data is changing or what is the forecast trend in the data are communicated. The PMCP system may transmit actual model parameters to receiving units such that edge devices can apply the vector-based predictive models to determine future states. For example, each automated device in a network may be configured to train a regression model or a neural network, constantly fitting the data streams to current input data. In example embodiments, automated devices leveraging the PMCP system are able to react in advance of events actually happening, rather than, for example, waiting for a machine malfunction to occur. Continuing the example, the stateless automated device can react to the forecast future state and make the necessary adjustments, such as ordering a new machine part.

[2589] In example embodiments, the PMCP system enables communicating vectorized information together with model parameters that allow predictive models on a receiving end to predict probabilities of future values. The vectorized information may be transmitted and processed to determine a number of probability-based states. For example, motion vectors and model parameters for predicting future locations based on motion vectors may be transmitted using PMCP, and a receiving location may use the motion vectors as inputs to a parameterized predictive model (e.g., a model that determines future locations of an item using the model parameters), which may generate probabilities that an item associated with a motion vector is in different locations. As another example, the PMCP system may support communicating vectorized sensor readings together with model parameters that allow current and/or future sensor readings to be predicted. Applied in an environment with large numbers of sensors with different accuracies and reliabilities, the probabilistic vector-based mechanism of the PMCP system allows large numbers, if not all, data streams to be used to produce refined models representing the current state, past states and likely future states of industrial entities (e.g., machines, services, and/or the like). Approximation methods may include importance sampling, and the resulting predictive model may be a particle filter, condensation algorithm, Monte Carlo localization, or other suitable model.

[2590] In example embodiments, the vector-based communication of the PMCP system allows devices and/or other systems to anticipate future security events. For example, a set of simple edge devices may be configured to run semi-autonomously using PMCP to generate and transmit model parameters based on locally-sensed security data. In this example, the edge devices may be configured to build a set of forecast models showing trends in the data. The parameters of this set of forecast models may be transmitted using the PMCP system. In this example, the edge devices may be configured to build a set of forecast models showing trends in the data. The parameters of this set of forecast models may be transmitted using the PMCP system so that the security data may be rebuilt and used to predict future states at a receiving device.

[2591] In example embodiments, security systems may generate and transmit vectors showing changes in state, as unusual events tend to cause one or more vectors to show unusual patterns. In a security setting, detecting multiple simultaneous unusual vectors may trigger escalation and a response by, for example, a control tower. In addition, one of the major areas of communication

security concern is around the protection of stored data, and in a vector-based system data may not need to be stored (or may be stored on fewer devices), so the risk of data loss is removed or reduced.

[2592] In example embodiments, PMCP data can be directly stored in a queryable database where the actual data is reconstructed dynamically in response to a query. In example embodiments, the PMCP data streams can be used to recreate the fine-grained data so they become part of an Extract Transform and Load (ETL) process.

[2593] A PMCP device interface may include several modules including a transceiver module 68062, an intelligence module 68064, a library module 68066, and a storage module 68068. The transceiver module may include a data transceiver 68070 that may be used to transmit/receive data, including various data from data sources 68004 and/or PMCP data (e.g., vectors, model parameters, etc.) to/from other PMCP device interfaces (e.g., PMCP device interface 68052) and/or to/from other components of a system including the PMCP device interface. In example embodiments, the transceiver module 68062 may include an intelligence system 68072, which may use artificial intelligence techniques to assist in transmission and/or reception processing. For example, the intelligence system 68072 may route various types of incoming and outgoing data, prioritize or deprioritize transmitted and/or received data from data sources 68004 vs PMCP data, and/or the like. The intake module 68074 may further include a PMCP controller 68074, which may understand PMCP transmissions, parse PMCP data, and provide the received PMCP data to the modelling module for further operations.

[2594] The modelling module 68064 may be responsible for various operations in a transmission role and/or in a receiver role. In a transmission role, the modelling module 68064 may continually receive data from various data sources 68004 (e.g., sensors 68022) and continually generate and/or refine models that predict future states of the incoming data. The various models may be, for example, classification models, behavioral analysis models, prediction models, data augmentation models, and/or any other types of model. Model parameters (e.g., neural network weights) from the generated/refined models may then be transmitted to receivers, which may use the parameters to perform classifications, behavior analysis, prediction, augmentation and/or the like without needing to have access to the data stream. Accordingly, in a receiver role, the modelling module 68064 may use various parameters received from another PMCP device interface to parameterize various types of models, then use the parameterized models to generate data for further use by the receiving device.

[2595] In example embodiments, the PMCP device interface may train and/or execute classification models 68076, which may be trained using data captured from data sources 68004 generate various labels or classifications. For example, classification models may be used to output various states or conditions based on input data, including predicted future states or conditions. By transmitting classification model parameters to a receiving device using PMCP, the receiving device may also be able to predict the future states or conditions without having to receive the input data from the data sources 68004.

[2596] In example embodiments, the PMCP device interface may train and/or execute behavior analysis models 68078, which may be trained using data captured from data sources 68004

generate various behavioral analyses and future behavioral data. For example, behavior analysis models may be used to output current or future actions that are likely to be taken by certain entities and/or analyses of whether the actions are within normal conditions or unusual. By transmitting behavioral analysis model parameters to a receiving device using PMCP, the receiving device may also be able to predict the future actions and/or analyses without having to receive the input data from the data sources 68004.

[2597] In example embodiments, the PMCP device interface may train and/or execute prediction models 68080, which may be trained using data streams captured from data sources 68004 generate current and predicted data values for the data streams. For example, prediction models may be used to output current or future sensor readings based on data captured from sensors 68022. By transmitting prediction model parameters to a receiving device using PMCP, the receiving device may also be able to predict the sensor values without having to receive the input data from the sensors 68002 or other data sources 68004.

[2598] In example embodiments, the PMCP device interface may train and/or execute augmentation models 68082, which may be trained using data captured from data sources 68004 to generate augmented data streams. For example, augmentation models may be used to generate interpolated or extrapolated values from data streams that may be missing data (e.g., due to network interruptions), may generate predicted sensor readings for a sensor (e.g., a broken sensor) based on sensor readings from other nearby sensors, and may otherwise augment data received from data sources 68004 with additional data. By transmitting augmentation model parameters to a receiving device using PMCP, the receiving device may also be able to generate the missing data, predicted data, or other augmented data without having to receive the input data from the data sources 68004.

[2599] In example embodiments, the PMCP device interface 68060 may use a library module 68066 containing one or more modules that may be used to assist in modelling and/or other operations. For example, a networking module 68084 may contain various data about network devices, network topologies, network digital twins, and other network data that may be leveraged to train various models, to perform ETL operations as described in more detail below, or to perform other such processing. As another example, a security module 68086 may contain various data about security devices, building layouts (e.g., for building security systems), maps, topologies, digital twins, vulnerabilities, and other security data that may be leveraged to train various models, to perform ETL operations as described in more detail below, or to perform other such processing for security reasons. Various other specific modules may be provided to enable or support specific use cases.

[2600] In example embodiments, a storage module 68068 may provide various operations for processing data for storage and/or storing data. An ETL interface 68088 may be configured to perform exchange, transform, and load (ETL) operations for storing data in a PMCP database 68090. The PMCP database 68090 may be used to store various data, including data received from data sources 68004 (e.g., such that historical data may be used to generate/refine various models), as well as the models themselves, model parameters, and/or the like.

[2601] In example embodiments, the thalamus service and PMCP may provide complementary techniques for managing large amounts of data. For example, PMCP may reduce the bandwidth and storage requirements for working with large amounts of data because PMCP may only require transmitting model parameters, instead of transmitting bandwidth-intensive data streams. However, when dealing with large numbers of data sensors or other data sources, PMCP may not be enough to reduce data to manageable levels, as the number of PMCP streams, number of models, etc. may still be too large to handle. In these cases, the thalamus service may operate to prioritize, format, suppress, filter, or combine PMCP data streams in order to allow for a focus on the most important PMCP data streams at any given time. Several benefits are realized by combining the techniques in this manner. For example, although massive amounts of data may be collected, PMCP may allow the communication of model parameters for predicting some or all of the data, and the thalamus service may allow for a focus on the most important models and predictions at any given time. Moreover, the use of PMCP causes the data to be inherently predictive and thus forward-looking, which, in combination with the thalamus service, allows for a focus on the most important data before the occurrence of potential issues that may need various actions (e.g., interventions, maintenance, purchase orders, supply adjustments, estimate adjustments, etc.).

[2602] FIG. 257 shows the interaction of the intake controller 68008, intake management system 68006, and various other components of the thalamus service 68000 with PMCP according to some embodiments of the invention. In the illustrated embodiments, inputs may be received to the intake controller 68008 from different sources. For example, a first source of data may include various sensors, external systems, process data, and other such data 18102 that may be received from various data generators, data analysis systems, and other data outputs outside of the thalamus service. Additionally or alternatively, a second source of data may include one or more preconfigured PMCP devices with location processing, which may provide at 18104 that may include PMCP model parameters, vectorized data, or other PMCP data.

[2603] The intake controller 68008 may ingest the data and determine whether the data is PMCP data or not at a decision 18106. If the data is not PMCP data, then the intake controller 68008 may determine if the data has been reduced or not. If the data has not been reduced, then the data may be sent to the intake management system for processing (e.g., prioritization, formatting, suppressing, area focus, filtering, combining, etc. as discussed above). In other words, if the data has not already been reduced in some way (e.g., either via PMCP or using other data reduction techniques), the data may be processing and potentially filtered, suppressed, or otherwise reduced. Thus, the thalamus service may provide data reduction techniques that may be used in addition to or as an alternative to other data reduction techniques, which may include PMCP.

[2604] If the data was not PMCP data but was reduced as determined at 18108, or if the data was PMCP data as determined at 68016, then the intake controller 18110 may determine whether the thalamus service is acting as a PMCP consumer for the data. If so, the data may be sent to the PMCP device interface 68052 for reception and processing (e.g., modelling, prediction, etc.). If not, then one or more ETL processes may be used at 18114 to extract, transform, and load the data into the PMCP database.

[2605] Whether the data is processed by the PMCP device interface 68052 or using ETL processes at 18114, the resulting data may then be provided to downstream data consumers for further processing at 18116.

[2606] PMCP and thalamus service techniques may be used (together or separately) in a wide variety of embodiments. In embodiments where edge devices are configured with very limited capacities, additional edge communication devices can be added to convert the data into PMCP format. For example, to protect distributed equipment from hacking attempts, many manufacturers will choose to not connect the device to any kind of network. To overcome this limitation, the equipment may be monitored using sensors, such as cameras, sound monitors, voltage detectors for power usage, chemical sniffers, and the like. Functional unit learning and other data techniques may be used to determine the actual usage of the equipment detached from the network functional unit, generate vectorized data therefrom, and/or transmit various model parameters using PMCP. On the receiving end, a thalamus service may receive the vectorized data and/or model parameters, may use thalamus techniques to determine whether the PMCP data and/or other data received from other equipment should be prioritized, filtered, suppressed, or the like, may predict future states of the equipment based on the PMCP data, and may use any or all of the data to take various actions, perform various analyses, and the like.

[2607] In example embodiments, communication using vectorized data allows for a constant view of what the likely future state is. These techniques allow for future states to be communicated, thus allowing industrial entities to respond ahead of future state requirements without needing access to fine-grained data.

[2608] In example embodiments, the PMCP protocol can be used to transmit and receive relevant information (e.g., important or high priority information, as determined by a thalamus service) about manufacturing performance indicators and future trends in manufacturing performance to various external entities. In some of these example embodiments, a PMCP data feed may be used for data obfuscation (e.g., communicating sensitive data as vectorized data and/or model parameters). For example, PMCP allows real contextual information about manufacturing performance to be shared with stakeholders, regulators, and other external entities without the direct sharing of sensitive data values.

[2609] PMCP and vectorized data processes further enable simple data-informed interactive systems that a user can apply without having to build enormously complex big data engines. As an example, an upstream manufacturer may have an enormously complex task of coordinating many downstream consumption points. Through the use of PMCP and/or thalamus services, the manufacturer may be able to provide real information to consumers without the need to store detailed data and build complex models, which may require setting up large-scale systems for processing large amounts of data and the like.

[2610] In example embodiments, edge device units may communicate via the PMCP system to show direction of movement and likely future positions. For example, a moving robot can communicate its likely track of future movement. In embodiments involving large numbers of moving robots, a thalamus service may determine which robots need to be prioritized and

monitored closely (e.g., because they are moving outside of prescribed boundaries, behaving in unpredictable ways, etc.).

[2611] In example embodiments, the PMCP system and/or thalamus system enables visual representations of vector-based data (e.g., via a user interface), including highlighting of areas of concern without the need to process enormous volumes of data. The visual representation allows for the display of many monitored vector inputs. The user interface can then display information relating to the key items of interest, specifically vectors showing areas of unusual or troublesome movement. This mechanism allows sophisticated models that are built at the edge device edge nodes to feed into end user communications in a visually informative way.

[2612] As can be appreciated, functional units produce a constant stream of "boring" data (e.g., data that does not change, changes slightly, or changes very predictably). By changing from producing data, to monitoring for problems, issues with the logistical modules are highlighted without the need for scrutiny of fine-grained data. In example embodiments, PMCP device interfaces may constantly generate and/or refine a predictive model that predicts a future state. In the context of maintenance, refinements to the parameters in the predictive model are in and of themselves predictors of change in operational parameters, potentially indicating the need for maintenance. Moreover, the communication of operational parameters for large numbers of devices may be processed by a thalamus service such that data for devices functioning normally may be filtered or suppressed until conditions change.

[2613] In example embodiments, functional areas are not always designed to be connected to a network, but by allowing for an external device to virtually monitor devices, functional areas that do not allow for connectivity can become part of the information flow. This concept extends to allowing functional areas that have limited connectivity to be monitored effectively by embellishing their data streams with vectorized monitored information. Placing an automated device in the proximity of the functional unit that has limited or no connectivity allows capture of information from the devices without the requirement of connectivity. There is also potential to add training data capture functional units for these unconnected or limitedly connected functional areas. These training data capture functional units are typically quite expensive and can provide high quality monitoring data, which is used as an input into the proximity edge device monitoring device to provide data for supervised learning algorithms.

[2614] Oftentimes, industrial locations are laden with electrical interference, causing fundamental challenges with communications. The traditional approach of streaming all the fine-grained data is dependent on the completeness of the data stream. For example, if an edge device were to go offline for 10 minutes, the streaming data and its information would be lost. With vectorized communication, the offline unit may continue to refine the predictive model until the moment when it reconnects, which allows the updated model to be transmitted via the PMCP system.

[2615] In example embodiments, industrial systems and devices may be based on the PMCP protocol. For example, industrial cameras and vision systems (e.g., liquid lens systems), user devices, sensors, robots, machines, and the like may use PMCP and/or vector-based communication. By using vector-based cameras, for example, only information relating to the

movement of items is transmitted. This reduces the data volume and by its nature filters information about static items, showing only the changes in the images and focusing the data communication on elements of change. The overall shift in communication to communication of change is similar to how the human process of sight functions, where stationary items are not even communicated to the higher levels of the brain.

[2616] Radio Frequency Identification allows for massive volumes of mobile to be tracked in real-time. In example embodiments, the movement of the tags may be communicated as vector information via the PMCP protocol, as this form of communication is naturally suited to handing information regarding the location of tag. Adding the ability to show future state of the location using predictive models that can use paths of prior movement allows the fundamental communication mechanism to be one in which units consuming data streams are consuming information about the likely future state of the industrial entities. In example embodiments, each tagged item may be represented as a probability-based location matrix showing the likely probability of the tagged item being at a position in space. The communication of movement shows the transformation of the location probability matrix to a new set of probabilities. This probabilistic locational overview provides for constant modeling of areas of likely intersection of moving units and allows for refinement of the probabilistic view of the location of items. Moving to a vector-based probability matrix allows units to constantly handle the inherent uncertainty in the measurement of status of industrial items, entities, and the like. In example embodiments, status includes, but is not limited to, location, temperature, movement and power consumption.

[2617] In example embodiments, continuous connectivity is not required for continuous monitoring of sensor inputs in a PMCP-based communication system. For example, a mobile robotic device with a plurality of sensors can continue to build models and predictions of data streams while disconnected from the network, and upon reconnection, the updated models are communicated. Furthermore, other systems or devices that use input from the monitored system or device can apply the best known, typically last communicated, vector predictions to continue to maintain a probabilistic understanding of the states of the industrial entities.

[2618] Techniques described herein improve the ability of networks and systems to collect, transmit, and process large volumes of data, especially data from sensors and other industrial internet of things data generators. These techniques include using a thalamus service that provides an equivalent to a biological thalamus, a neural system for filtering and relaying data. The thalamus service described herein can receive large volumes of information and quickly prioritize the information, passing on the most importing information so that limited transmission, processing, collection, and/or analysis resources are not overwhelmed by volume of incoming information.

[2619] Additionally, a predictive model communication protocol (PMCP) is described herein. PMCP may be used to reduce a volume of transmitted data, especially when the data is predictable or usually predictable. PMCP may operate by transmitting predictive model parameters instead of some or all of the data values that would normally be transmitted by a sensor device or other data source. For example, a device implementing PMCP may continually receive inputs (e.g., sensor data) and train a predictive model using the stream of sensor data. Rather than transmitting the

sensor data, which may use significant network and/or processing resources, the PMCP device may transmit the model parameters, which may be used by a receiving device to operate a predictive model to predict current and future sensor data. Thus, the receiving device may have a predictive model of sensor data without receiving the sensor data. In example embodiments, if the sensor data at the PMCP device begins operating outside of expectations, the model parameters may be re-transmitted to the receiving device, which may update its predictive model and thereby obtain more accurate predictive data.

[2620] In example embodiments, to optimize decision-making, quantum computers and/or predictive models may be used with the techniques described herein. Furthermore, quantum coordination can be applied to allow for disparate units to securely coordinate actions (e.g., without the need for traditional communication mechanisms). Accordingly, techniques described herein may use a combination of decentralized biology-based decision-making capabilities distributed throughout devices within the industrial internet of things and quantum capabilities.

[2621] FIG. 256 shows a thalamus service 68000 and a set of input sensors streaming data from various sources across the industrial internet of things and the control system 68002 with its centrally-managed data sources 68004. The thalamus service 68000 filters the inputs from the various data sources 68004 into the control system 68002 such that the control system is never overwhelmed by the total volume of information. In example embodiments, the thalamus service 68000 provides an information suppression mechanism for information flows. This mechanism monitors all data streams and suppresses and/or filters irrelevant data streams by ensuring that the maximum data flows from all input sensors are always constrained.

[2622] In example embodiments, the thalamus service 68000 may be a gateway for all communication that responds to the prioritization of the control system 68002. The control system 68002 may decide to change the prioritization of the data streamed from the thalamus service 68000, for example, during a known fire in an isolated area, and the event may direct the thalamus service 68000 to continue to provide flame sensor information despite the fact that majority of this data is not unusual. The thalamus service 68000 may be an integral part of the overall communication framework.

[2623] In example embodiments, the thalamus service 68000 includes an intake management system 68006. The intake management system 68006 may be configured to receive and process multiple large datasets by converting them into data streams that are sized and organized for subsequent use by a central control system 68002. For example, a robot may include vision and sensing systems that are used by the central control system 68002 (which may be on-board the robot and/or in a separate device in communication with the robot) to identify and move through an environment in real-time. The intake management system 68006 can facilitate robot decision-making by parsing, filtering, classifying, or otherwise reducing the size and increasing the utility of multiple large datasets that would otherwise overwhelm the central control system 68002. In example embodiments, the intake management system may include an intake controller 68008 that works with the intelligence service 68010 to evaluate incoming data and take actions-based evaluation results. Evaluations and actions may include specific instruction sets received by the

thalamus service 68000, for example the use of a set of specific compression and prioritization tools stipulated within a "Networking" library module. In another example, thalamus service inputs may direct the use of specific filtering and suppression techniques. In a third example, thalamus service inputs may stipulate data filtering associated with an area of interest such as a certain type of machine. The intake management system is also configured to recognize and manage datasets that are in a vectorized format such as in accordance with a predictive model communication protocol (PMCP) (discussed below), where the datasets may be passed directly to the central control system 68002, or alternatively deconstructed and processed separately. The intake management system 68006 may include a learning module that receives data from external sources that enables improvement and creation of application and data management library modules. In some cases, the intake management system 68006 may request external data to augment existing datasets.

[2624] In example embodiments, the control system 68002 may direct the thalamus service 68000 to alter its filtering to provide more input from a set of specific sources. This indication to provide more input is handled by the thalamus service 68000. For example, the thalamus service may suppress other information flows to constrain the total data flows to within a volume that the central control system can handle.

[2625] In example embodiments, the thalamus service 68000 can operate by suppressing data based on several different factors including zero or more default factors. For example, in some embodiments the default factors may include an "unusualness factor" that may be a value that indicates a divergence or a degree of divergence of the data from an expected dataset. In example embodiments, the unusualness factor is constantly monitored for all inputs or some of the inputs (e.g., some of the input sensors).

[2626] In example embodiments, the thalamus service 68000 may suppress data based on geospatial factors. Examples of geospatial factors may include location data, motion data, acceleration data, vibration data, and/or any other data indicating an absolute or relative location, change in location over time, other derivatives or integrals of location over time, etc. The thalamus service 68000 may be aware of the geospatial factors for some or all of the sensors and thus is able to look for unusual patterns in data based on geospatial context and suppress data accordingly.

[2627] In example embodiments, the thalamus service 68000 may suppress data based on temporal factors. Data can be suppressed temporally, for example, if the cadence of the data can be reduced such that the overall data stream is filtered to a level that can be handled by the control system 68002 and/or a central processing unit.

[2628] In example embodiments, the thalamus service 68000 may suppress data based on contextual factors. In example embodiments, context-based filtering is a filtering event in which the thalamus service 68000 is aware of some context-based event. Context-based events, for example, may include one or more notifications of unusual behavior by other sensors or systems (which may lead to temporary suppression of less important data), one or more human inputs (e.g., a human disabling a security alert, which may suppress a previous focus on security data), one or more events triggered by other systems or sensors (e.g., an automated security alert, which may

lead to suppression of certain data to allow resources to be dedicated to security data collection, transmission, and analysis), one or more contexts detected from other sensor data (e.g., a reduction in available bandwidth reported by a network sensor, which may lead to the suppression of certain data until available bandwidth improves), or any other context-based condition or event. In this context, the filtering may suppress information flows not relating to the data from the event.

[2629] In example embodiments, the thalamus service 68000 may receive data from a variety of data sources 68004, including analyses 68018, databases 68020, sensors 68022, and/or reports 68024. For example, the thalamus service 68000 may receive analyses and/or reports from other analysis/processing/reporting devices that have already pre-processed sensor data or other data. Additionally or alternatively, the thalamus service 68000 may receive data (e.g., historical data) that is stored in a database 68020 in addition to current or historical data from sensors 68022. In example embodiments, data may be received and/or generated (e.g., predictive models may generate future data) from the PMCP device interface 68052.

[2630] In example embodiments, the thalamus service 68000 may process and/or interpret inputs from any of the data sources 68004 based on an intake application library 68012, which may include a networking library 68014, a security library 68016, and/or any other library for interpreting various types of input data. For example, the thalamus service 68000 may use a networking 68014 library to parse, interpret, extract, and/or otherwise process network data (e.g., data received from networking sensors or devices, networking analyses, networking reports, network database data, etc.). Similarly, the thalamus service 68000 may use a security 68016 library to parse, interpret, extract, and/or otherwise process security data (e.g., data received from security sensors or devices, security analyses, security reports, security database data, etc.). In example embodiments, the intake data may also be processed using an intake learning module 68026, which may use one or more artificial intelligence techniques to pre-process the data, generate predictive models using the data, predict future states of the data, and/or the like. After processing using the intake application library 68012 and/or the intake learning module 68026, the data may be ready for management by the intake data management system 68028.

[2631] The intake data management system 68028 may process the data by prioritizing 68030, formatting 68032, suppressing 68034, using an area focus 68036, filtering 68038, and/or combining 68040 the data. The prioritizing 68030 may involve ranking or otherwise assigning priorities (e.g., categories, numerical priority scores, etc.) such that limited resources may be assigned to the most important data. For example, the suppressing 68034 and/or filtering 68038 may operate based on priorities in order to suppress or filter out the least important data (e.g., the data associated with a lowest priority score) in order to avoid overwhelming limited transmission, processing, and/or analysis resources. The formatting 68032 may involve formatting data in order to allow for easier management, which may involve compressing or otherwise dropping certain parts of data to reduce use of transmission resources, un-compressing data to reduce use of decompression resources (e.g., if bandwidth is sufficient and data is important), formatting data to emphasize or de-emphasize certain aspects, or otherwise adjust formatting. In example embodiments, the formatting 68032 may depend on the prioritizing 68030 such that more

important data may be formatted in order to allow for more or better analysis, while less important data may be formatted in order to reduce it usage of various resources.

[2632] In example embodiments, the suppressing 68034 may involve reducing the amount of data, the number of destinations to which the data is transmitted, or otherwise reducing the usage of limited resources (e.g., bandwidth, processing, analysis, etc.) of the data. In example embodiments, suppressed data may be stored (e.g., in a database) and dealt with (e.g., transmitted, processed) at a later time. In example embodiments, the suppressing 68034 may be based various factors as described above.

[2633] In example embodiments, an area focus 68036 may involve increasing the attention paid to certain high priority data. For example, during a security incident, security sensor data may be sent to additional destinations, processed using additional analyses, allowed additional bandwidth and processing power, and/or the like. In example embodiments, an area focus 68036 may cause the suppression or filtering of other data that is not associated with the area focus 68036.

[2634] In example embodiments, the filtering 68038 may involve ignoring, deleting, or otherwise removing data that is not important (e.g., does not match an area focus 68036, is low priority, etc.). In example embodiments, data may be initially suppressed (e.g., reduced or stored for later), but conditions may further change, causing the data to be filtered (e.g., deleted, ignored). Thus, intake data management system 68028 may allow for a progressive downgrade of data by first suppressing and later filtering the data depending on conditions.

[2635] In example embodiments, the combining 68040 may include combining various types of data in order to provide better analyses, generate new data, reduce the volume of data (e.g., by combining multiple data values into a single data value), improve the quality of data (e.g., by averaging different sensor readings to obtain a more accurate average reading), and/or the like. In example embodiments, lower priority data may be combined with other data in order to reduce requirements. Additionally or alternatively, higher priority data may be combined with other data in order to improve data quality.

[2636] In example embodiments, the intake data management system may interface with an intake controller 68008 and/or an intelligence system 68042. The intelligence system 68042, for example, may use various artificial intelligence techniques to perform the intake data management (e.g., prioritize the data, format the data, suppress the data, select an area focus and/or assign data to an area focus, filter the data, combine the data, etc.), predict the outcomes of intake data management, predict future data values, and/or the like. Additionally or alternatively, the intake controller 68008 and/or an intelligence system 68042 may operate in accordance with configured thalamus parameters 68044, which may govern the intake data management, the artificial intelligence techniques (e.g., the parameters may be model parameters for AI models), and/or otherwise configure the operations of the intake management system 68006.

[2637] In example embodiments, the control system 68002 may, in some cases, use a quantum computing service 68046, which may provide quantum computing resources to more quickly process large volumes of data, use quantum models, and/or the like.

[2638] The control system 68002 may further comprise one or more data interfaces 68048 for receiving data from various data sources 68004 and transmitting the data (e.g., after intake data management) to various destinations. In example embodiments, the control system 68002 may include other system subsystems 68050, such as analysis subsystems, various processing chips, or any other subsystems that may use the managed data to make decisions, generate analyses, or otherwise perform data operations. In example embodiments, an intelligence service 68010 may operate to route the managed data to various other system subsystems 68050, or otherwise perform initial and/or final processing on the data.

[2639] In example embodiments, the control system 68002 can override the thalamus filtering and decide to focus on a different area for any specific reason. For example, during a security incident, the control system may route around thalamus filtering (which might normally de-prioritize data from security sensors) in order to ensure that data from security sensors is delivered in full without any de-prioritization, suppression, filtering, etc. As another example, during regular inspections of equipment, sensor data that measures operation of the equipment (e.g., vibration sensor data) may be un-suppressed, even if the data appears to be within normal parameters and therefore might usually be suppressed or filtered.

[2640] In example embodiments, the control system 68002 may include a PMCP device interface 68052, which may be used to transmit and/or receive data using PMCP. Details of a PMCP device interface are further shown within a second PMCP device interface 68060. In example embodiments, the PMCP device interface 68052 may be in communication with the PMCP device interface 68060. The PMCP device interface 68052 may have the same components as shown within the PMCP device interface 68060.

[2641] In example embodiments, the PMCP device interface may be used to convert data to a vectorized format prior to transmission. In these embodiments, a vector may be considered an example of a simple predictive model (e.g., a vector may indicate an amount of change and a direction of change for a data value, thus predicting a future state of a data value if the change continues). For example, the conversion of a long sequence of oftentimes similar data values into a vector indicating an amount and direction of change, which may imply a future state of the data values, makes the communication of the data values both smaller in size and forward looking in nature.

[2642] In example embodiments, PMCP may use various types of predictive models to predict current and future data values, including weighted moving average; Kalman filtering; exponential smoothing; autoregressive moving average (ARMA) (forecasts depend on past values of the variable being forecast, and on past prediction errors); autoregressive integrated moving average (ARIMA) (ARMA on the period-to-period change in the forecasted variable); extrapolation; linear prediction; trend estimation (predicting the variable as a linear or polynomial function of time); growth curve (e.g., statistics); and recurrent neural network based forecasting.

[2643] Using the PMCP protocol, instead of traditional streams where individual data items are transmitted, vectors representing how the data is changing or what is the forecast trend in the data are communicated. The PMCP system may transmit actual model parameters to receiving units

such that edge devices can apply the vector-based predictive models to determine future states. For example, each automated device in a network may be configured to train a regression model or a neural network, constantly fitting the data streams to current input data. In example embodiments, automated devices leveraging the PMCP system are able to react in advance of events actually happening, rather than, for example, waiting for a machine malfunction to occur. Continuing the example, the stateless automated device can react to the forecast future state and make the necessary adjustments, such as ordering a new machine part.

[2644] In example embodiments, the PMCP system enables communicating vectorized information together with model parameters that allow predictive models on a receiving end to predict probabilities of future values. The vectorized information may be transmitted and processed to determine a number of probability-based states. For example, motion vectors and model parameters for predicting future locations based on motion vectors may be transmitted using PMCP, and a receiving location may use the motion vectors as inputs to a parameterized predictive model (e.g., a model that determines future locations of an item using the model parameters), which may generate probabilities that an item associated with a motion vector is in different locations. As another example, the PMCP system may support communicating vectorized sensor readings together with model parameters that allow current and/or future sensor readings to be predicted. Applied in an environment with large numbers of sensors with different accuracies and reliabilities, the probabilistic vector-based mechanism of the PMCP system allows large numbers, if not all, data streams to be used to produce refined models representing the current state, past states and likely future states of industrial entities (e.g., machines, services, and/or the like). Approximation methods may include importance sampling, and the resulting predictive model may be a particle filter, condensation algorithm, Monte Carlo localization, or other suitable model.

[2645] In example embodiments, the vector-based communication of the PMCP system allows devices and/or other systems to anticipate future security events. For example, a set of simple edge devices may be configured to run semi-autonomously using PMCP to generate and transmit model parameters based on locally-sensed security data. In this example, the edge devices may be configured to build a set of forecast models showing trends in the data. The parameters of this set of forecast models may be transmitted using the PMCP system. In this example, the edge devices may be configured to build a set of forecast models showing trends in the data. The parameters of this set of forecast models may be transmitted using the PMCP system so that the security data may be rebuilt and used to predict future states at a receiving device.

[2646] In example embodiments, security systems may generate and transmit vectors showing changes in state, as unusual events tend to cause one or more vectors to show unusual patterns. In a security setting, detecting multiple simultaneous unusual vectors may trigger escalation and a response by, for example, a control tower. In addition, one of the major areas of communication security concern is around the protection of stored data, and in a vector-based system data may not need to be stored (or may be stored on fewer devices), so the risk of data loss is removed or reduced.

[2647] In example embodiments, PMCP data can be directly stored in a queryable database where the actual data is reconstructed dynamically in response to a query. In example embodiments, the

PMCP data streams can be used to recreate the fine-grained data so they become part of an Extract Transform and Load (ETL) process.

[2648] A PMCP device interface may include several modules including a transceiver module 68062, an intelligence module 68064, a library module 68066, and a storage module 68068. The transceiver module may include a data transceiver 68070 that may be used to transmit/receive data, including various data from data sources 68004 and/or PMCP data (e.g., vectors, model parameters, etc.) to/from other PMCP device interfaces (e.g., PMCP device interface 68052) and/or to/from other components of a system including the PMCP device interface. In example embodiments, the transceiver module 68062 may include an intelligence system 68072, which may use artificial intelligence techniques to assist in transmission and/or reception processing. For example, the intelligence system 68072 may route various types of incoming and outgoing data, prioritize or deprioritize transmitted and/or received data from data sources 68004 vs PMCP data, and/or the like. The intake module 68074 may further include a PMCP controller 68074, which may understand PMCP transmissions, parse PMCP data, and provide the received PMCP data to the modelling module for further operations.

[2649] The modelling module 68064 may be responsible for various operations in a transmission role and/or in a receiver role. In a transmission role, the modelling module 68064 may continually receive data from various data sources 68004 (e.g., sensors 68022) and continually generate and/or refine models that predict future states of the incoming data. The various models may be, for example, classification models, behavioral analysis models, prediction models, data augmentation models, and/or any other types of model. Model parameters (e.g., neural network weights) from the generated/refined models may then be transmitted to receivers, which may use the parameters to perform classifications, behavior analysis, prediction, augmentation and/or the like without needing to have access to the data stream. Accordingly, in a receiver role, the modelling module 68064 may use various parameters received from another PMCP device interface to parameterize various types of models, then use the parameterized models to generate data for further use by the receiving device.

[2650] In example embodiments, the PMCP device interface may train and/or execute classification models 68076, which may be trained using data captured from data sources 68004 generate various labels or classifications. For example, classification models may be used to output various states or conditions based on input data, including predicted future states or conditions. By transmitting classification model parameters to a receiving device using PMCP, the receiving device may also be able to predict the future states or conditions without having to receive the input data from the data sources 68004.

[2651] In example embodiments, the PMCP device interface may train and/or execute behavior analysis models 68078, which may be trained using data captured from data sources 68004 generate various behavioral analyses and future behavioral data. For example, behavior analysis models may be used to output current or future actions that are likely to be taken by certain entities and/or analyses of whether the actions are within normal conditions or unusual. By transmitting behavioral analysis model parameters to a receiving device using PMCP, the receiving device may

also be able to predict the future actions and/or analyses without having to receive the input data from the data sources 68004.

[2652] In example embodiments, the PMCP device interface may train and/or execute prediction models 68080, which may be trained using data streams captured from data sources 68004 generate current and predicted data values for the data streams. For example, prediction models may be used to output current or future sensor readings based on data captured from sensors 68022. By transmitting prediction model parameters to a receiving device using PMCP, the receiving device may also be able to predict the sensor values without having to receive the input data from the sensors 68002 or other data sources 68004.

[2653] In example embodiments, the PMCP device interface may train and/or execute augmentation models 68082, which may be trained using data captured from data sources 68004 to generate augmented data streams. For example, augmentation models may be used to generate interpolated or extrapolated values from data streams that may be missing data (e.g., due to network interruptions), may generate predicted sensor readings for a sensor (e.g., a broken sensor) based on sensor readings from other nearby sensors, and may otherwise augment data received from data sources 68004 with additional data. By transmitting augmentation model parameters to a receiving device using PMCP, the receiving device may also be able to generate the missing data, predicted data, or other augmented data without having to receive the input data from the data sources 68004.

[2654] In example embodiments, the PMCP device interface 68060 may use a library module 68066 containing one or more modules that may be used to assist in modelling and/or other operations. For example, a networking module 68084 may contain various data about network devices, network topologies, network digital twins, and other network data that may be leveraged to train various models, to perform ETL operations as described in more detail below, or to perform other such processing. As another example, a security module 68086 may contain various data about security devices, building layouts (e.g., for building security systems), maps, topologies, digital twins, vulnerabilities, and other security data that may be leveraged to train various models, to perform ETL operations as described in more detail below, or to perform other such processing for security reasons. Various other specific modules may be provided to enable or support specific use cases.

[2655] In example embodiments, a storage module 68068 may provide various operations for processing data for storage and/or storing data. An ETL interface 68088 may be configured to perform exchange, transform, and load (ETL) operations for storing data in a PMCP database 68090. The PMCP database 68090 may be used to store various data, including data received from data sources 68004 (e.g., such that historical data may be used to generate/refine various models), as well as the models themselves, model parameters, and/or the like.

[2656] In example embodiments, the thalamus service and PMCP may provide complementary techniques for managing large amounts of data. For example, PMCP may reduce the bandwidth and storage requirements for working with large amounts of data because PMCP may only require transmitting model parameters, instead of transmitting bandwidth-intensive data streams. However, when dealing with large numbers of data sensors or other data sources, PMCP may not

be enough to reduce data to manageable levels, as the number of PMCP streams, number of models, etc. may still be too large to handle. In these cases, the thalamus service may operate to prioritize, format, suppress, filter, or combine PMCP data streams in order to allow for a focus on the most important PMCP data streams at any given time. Several benefits are realized by combining the techniques in this manner. For example, although massive amounts of data may be collected, PMCP may allow the communication of model parameters for predicting some or all of the data, and the thalamus service may allow for a focus on the most important models and predictions at any given time. Moreover, the use of PMCP causes the data to be inherently predictive and thus forward-looking, which, in combination with the thalamus service, allows for a focus on the most important data before the occurrence of potential issues that may need various actions (e.g., interventions, maintenance, purchase orders, supply adjustments, estimate adjustments, etc.).

[2657] FIG. 257 shows the interaction of the intake controller 68008, intake management system 68006, and various other components of the thalamus service 68000 with PMCP according to some embodiments of the invention. In the illustrated embodiments, inputs may be received to the intake controller 68008 from different sources. For example, a first source of data may include various sensors, external systems, process data, and other such data 18102 that may be received from various data generators, data analysis systems, and other data outputs outside of the thalamus service. Additionally or alternatively, a second source of data may include one or more preconfigured PMCP devices with location processing, which may provide at 18104 that may include PMCP model parameters, vectorized data, or other PMCP data.

[2658] The intake controller 68008 may ingest the data and determine whether the data is PMCP data or not at a decision 18106. If the data is not PMCP data, then the intake controller 68008 may determine if the data has been reduced or not. If the data has not been reduced, then the data may be sent to the intake management system for processing (e.g., prioritization, formatting, suppressing, area focus, filtering, combining, etc. as discussed above). In other words, if the data has not already been reduced in some way (e.g., either via PMCP or using other data reduction techniques), the data may be processed and potentially filtered, suppressed, or otherwise reduced. Thus, the thalamus service may provide data reduction techniques that may be used in addition to or as an alternative to other data reduction techniques, which may include PMCP.

[2659] If the data was not PMCP data but was reduced as determined at 18108, or if the data was PMCP data as determined at 68016, then the intake controller 18110 may determine whether the thalamus service is acting as a PMCP consumer for the data. If so, the data may be sent to the PMCP device interface 68052 for reception and processing (e.g., modelling, prediction, etc.). If not, then one or more ETL processes may be used at 18114 to extract, transform, and load the data into the PMCP database.

[2660] Whether the data is processed by the PMCP device interface 68052 or using ETL processes at 18114, the resulting data may then be provided to downstream data consumers for further processing at 18116.

[2661] PMCP and thalamus service techniques may be used (together or separately) in a wide variety of embodiments. In embodiments where edge devices are configured with very limited

capacities, additional edge communication devices can be added to convert the data into PMCP format. For example, to protect distributed equipment from hacking attempts, many manufacturers will choose to not connect the device to any kind of network. To overcome this limitation, the equipment may be monitored using sensors, such as cameras, sound monitors, voltage detectors for power usage, chemical sniffers, and the like. Functional unit learning and other data techniques may be used to determine the actual usage of the equipment detached from the network functional unit, generate vectorized data therefrom, and/or transmit various model parameters using PMCP. On the receiving end, a thalamus service may receive the vectorized data and/or model parameters, may use thalamus techniques to determine whether the PMCP data and/or other data received from other equipment should be prioritized, filtered, suppressed, or the like, may predict future states of the equipment based on the PMCP data, and may use any or all of the data to take various actions, perform various analyses, and the like.

[2662] In example embodiments, communication using vectorized data allows for a constant view of what the likely future state is. These techniques allow for future states to be communicated, thus allowing industrial entities to respond ahead of future state requirements without needing access to fine-grained data.

[2663] In example embodiments, the PMCP protocol can be used to transmit and receive relevant information (e.g., important or high priority information, as determined by a thalamus service) about manufacturing performance indicators and future trends in manufacturing performance to various external entities. In some of these example embodiments, a PMCP data feed may be used for data obfuscation (e.g., communicating sensitive data as vectorized data and/or model parameters). For example, PMCP allows real contextual information about manufacturing performance to be shared with stakeholders, regulators, and other external entities without the direct sharing of sensitive data values.

[2664] PMCP and vectorized data processes further enable simple data-informed interactive systems that a user can apply without having to build enormously complex big data engines. As an example, an upstream manufacturer may have an enormously complex task of coordinating many downstream consumption points. Through the use of PMCP and/or thalamus services, the manufacturer may be able to provide real information to consumers without the need to store detailed data and build complex models, which may require setting up large-scale systems for processing large amounts of data and the like.

[2665] In example embodiments, edge device units may communicate via the PMCP system to show direction of movement and likely future positions. For example, a moving robot can communicate its likely track of future movement. In embodiments involving large numbers of moving robots, a thalamus service may determine which robots need to be prioritized and monitored closely (e.g., because they are moving outside of prescribed boundaries, behaving in unpredictable ways, etc.).

[2666] In example embodiments, the PMCP system and/or thalamus system enables visual representations of vector-based data (e.g., via a user interface), including highlighting of areas of concern without the need to process enormous volumes of data. The visual representation allows

for the display of many monitored vector inputs. The user interface can then display information relating to the key items of interest, specifically vectors showing areas of unusual or troublesome movement. This mechanism allows sophisticated models that are built at the edge device edge nodes to feed into end user communications in a visually informative way.

[2667] As can be appreciated, functional units produce a constant stream of "boring" data (e.g., data that does not change, changes slightly, or changes very predictably). By changing from producing data, to monitoring for problems, issues with the logistical modules are highlighted without the need for scrutiny of fine-grained data. In example embodiments, PMCP device interfaces may constantly generate and/or refine a predictive model that predicts a future state. In the context of maintenance, refinements to the parameters in the predictive model are in and of themselves predictors of change in operational parameters, potentially indicating the need for maintenance. Moreover, the communication of operational parameters for large numbers of devices may be processed by a thalamus service such that data for devices functioning normally may be filtered or suppressed until conditions change.

[2668] In example embodiments, functional areas are not always designed to be connected to a network, but by allowing for an external device to virtually monitor devices, functional areas that do not allow for connectivity can become part of the information flow. This concept extends to allowing functional areas that have limited connectivity to be monitored effectively by embellishing their data streams with vectorized monitored information. Placing an automated device in the proximity of the functional unit that has limited or no connectivity allows capture of information from the devices without the requirement of connectivity. There is also potential to add training data capture functional units for these unconnected or limitedly connected functional areas. These training data capture functional units are typically quite expensive and can provide high quality monitoring data, which is used as an input into the proximity edge device monitoring device to provide data for supervised learning algorithms.

[2669] Oftentimes, industrial locations are laden with electrical interference, causing fundamental challenges with communications. The traditional approach of streaming all the fine-grained data is dependent on the completeness of the data stream. For example, if an edge device were to go offline for 10 minutes, the streaming data and its information would be lost. With vectorized communication, the offline unit may continue to refine the predictive model until the moment when it reconnects, which allows the updated model to be transmitted via the PMCP system.

[2670] In example embodiments, industrial systems and devices may be based on the PMCP protocol. For example, industrial cameras and vision systems (e.g., liquid lens systems), user devices, sensors, robots, machines, and the like may use PMCP and/or vector-based communication. By using vector-based cameras, for example, only information relating to the movement of items is transmitted. This reduces the data volume and by its nature filters information about static items, showing only the changes in the images and focusing the data communication on elements of change. The overall shift in communication to communication of change is similar to how the human process of sight functions, where stationary items are not even communicated to the higher levels of the brain.

[2671] Radio Frequency Identification allows for massive volumes of mobile to be tracked in real-time. In example embodiments, the movement of the tags may be communicated as vector information via the PMCP protocol, as this form of communication is naturally suited to handing information regarding the location of tag. Adding the ability to show future state of the location using predictive models that can use paths of prior movement allows the fundamental communication mechanism to be one in which units consuming data streams are consuming information about the likely future state of the industrial entities. In example embodiments, each tagged item may be represented as a probability-based location matrix showing the likely probability of the tagged item being at a position in space. The communication of movement shows the transformation of the location probability matrix to a new set of probabilities. This probabilistic locational overview provides for constant modeling of areas of likely intersection of moving units and allows for refinement of the probabilistic view of the location of items. Moving to a vector-based probability matrix allows units to constantly handle the inherent uncertainty in the measurement of status of industrial items, entities, and the like. In example embodiments, status includes, but is not limited to, location, temperature, movement and power consumption.

[2672] In example embodiments, continuous connectivity is not required for continuous monitoring of sensor inputs in a PMCP-based communication system. For example, a mobile robotic device with a plurality of sensors can continue to build models and predictions of data streams while disconnected from the network, and upon reconnection, the updated models are communicated. Furthermore, other systems or devices that use input from the monitored system or device can apply the best known, typically last communicated, vector predictions to continue to maintain a probabilistic understanding of the states of the industrial entities.

[2673] In example embodiments, the platform 100 includes a dual process artificial neural network (DPANN) system 70000, as shown in FIG. 258. The DPANN system 70000 includes an artificial neural network (ANN) having behaviors and operational processes (such as decision-making) that are products of a training system and a retraining system. The training system is configured to perform automatic, trained execution of ANN operations. The retraining system performs effortful, analytical, intentional retraining of the ANN, such as based on one or more relevant aspects of the ANN, such as memory, one or more input data sets (including time information with respect to elements in such data sets), one or more goals or objectives (including ones that may vary dynamically, such as periodically and/or based on contextual changes, such as ones relating to the usage context of the ANN), and/or others. In cases involving memory-based retraining, the memory may include original/historical training data and refined training data. The DPANN system 70000 includes a dual process learning function (DPLF) configured to manage and perform an ongoing data retention process. The DPLF (including, where applicable, memory management process) facilitate retraining and refining of behavior of the ANN. The DPLF provides a framework by which the ANN creates outputs such as predictions, classifications, recommendations, conclusions and/or other outputs based on a historic inputs, new inputs, and new outputs (including outputs configured for specific use cases, including ones determined by parameters of the context of utilization (which may include performance parameters such as latency parameters, accuracy

parameters, consistency parameters, bandwidth utilization parameters, processing capacity utilization parameters, prioritization parameters, energy utilization parameters, and many others).

[2674] In example embodiments, the DPANN system 70000 stores training data, thereby allowing for constant retraining based on results of decisions, predictions, and/or other operations of the ANN, as well as allowing for analysis of training data upon the outputs of the ANN. The management of entities stored in the memory allows the construction and execution of new models, such as ones that may be processed, executed or otherwise performed by or under management of the training system. The DPANN system 70000 uses instances of the memory to validate actions (e.g., in a manner similar to the thinking of a biological neural network (including retrospective or self-reflective thinking about whether actions that were undertaken under a given situation where optimal) and perform training of the ANN, including training that intentionally feeds the ANN with appropriate sets of memories (i.e., ones that produce favorable outcomes given the performance requirements for the ANN).

[2675] In example embodiments, the DPLF may be or include the continued process retention of one or more training datasets and/or memories stored in the memory over time. The DPLF thereby allows the ANN to apply existing neural functions and draw upon sets of past events (including ones that are intentionally varied and/or curated for distinct purposes), such as to frame understanding of and behavior within present, recent, and/or new scenarios, including in simulations, during training processes, and in fully operational deployments of the ANN. The DPLF may provide the ANN with a framework by which the ANN may analyze, evaluate, and/or manage data, such as data related to the past, present and future. As such, the DPLF plays a crucial role in training and retraining the ANN via the training system and the retraining system.

[2676] In example embodiments, the DPLF is configured to perform a dual-process operation to manage existing training processes and is also configured to manage and/or perform new training processes, i.e., retraining processes. In example embodiments, each instance of the ANN is trained via the training system and configured to be retrained via the retraining system. The ANN encodes training and/or retraining datasets, stores the datasets, and retrieves the datasets during both training via the training system and retraining via the retraining system. The DPANN system 70000 may recognize whether a dataset (the term dataset in this context optionally including various subsets, supersets, combinations, permutations, elements, metadata, augmentations, or the like, relative to a base dataset used for training or retraining), storage activity, processing operation and/or output, has characteristics that natively favor the training system versus the retraining system based on its respective inputs, processing (e.g., based on its structure, type, models, operations, execution environment, resource utilization, or the like) and/or outcomes (including outcome types, performance requirements (including contextual or dynamic requirements), and the like. For example, the DPANN system 70000 may determine that poor performance of the training system on a classification task may indicate a novel problem for which the training of the ANN was not adequate (e.g., in type of data set, nature of input models and/or feedback, quantity of training data, quality of tagging or labeling, quality of supervision, or the like), for which the processing operations of the ANN are not well-suited (e.g., where they are prone to known

vulnerabilities due to the type of neural network used, the type of models used, etc.), and that may be solved by engaging the retraining system to retrain the model to teach the model to learn to solve the new classification problem (e.g., by feeding it many more labeled instances of correctly classified items). With periodic or continuous evaluation of the performance of the ANN, the DPANN system may subsequently determine that highly stable performance of the ANN (such as where only small improvements of the ANN occur over many iterations of retraining by the retraining system) indicates readiness for the training system to replace the retraining system (or be weighted more favorably where both are involved). Over longer periods of time, cycles of varying performance may emerge, such as where a series of novel problems emerge, such that the retraining system of the DPANN is serially engaged, as needed, to retrain the ANN and/or to augment the ANN by providing a second source of outputs (which may be fused or combined with ANN outputs to provide a single result (with various weightings across them), or may be provided in parallel, such as enabling comparison, selection, averaging, or context- or situation-specific application of the respective outputs).

[2677] In example embodiments, the ANN is configured to learn new functions in conjunction with the collection of data according to the dual-process training of the ANN via the training system and the retraining system. The DPANN system 70000 performs analysis of the ANN via the training system and performs initial training of the ANN such that the ANN gains new internal functions (or internal functions are subtracted or modified, such as where existing functions are not contributing to favorable outcomes). After the initial training, the DPANN system 70000 performs retraining of the ANN via the retraining system. To perform the retraining, the retraining system evaluates the memory and historic processing of the ANN to construct targeted DPLF processes for retraining. The DPLF processes may be specific to identified scenarios. The ANN processes can run in parallel with the DPLF processes. By way of example, the ANN may function to operate a particular make and model of a self-driving car after the initial training by the training system. The DPANN system 70000 may perform retraining of the functions of the ANN via the retraining system, such as to allow the ANN to operate a different make and model of car (such as one with different cameras, accelerometers and other sensors, different physical characteristics, different performance requirements, and the like), or even a different kind of vehicle, such as a bicycle or a spaceship.

[2678] In example embodiments, as quality of outputs and/or operations of the ANN improves, and as long as the performance requirements and the context of utilization for the ANN remain fairly stable, performing the dual-process training process can become a decreasingly demanding process. As such, the DPANN system 70000 may determine that fewer neurons of the ANN are required to perform operations and/or processes of the ANN, that performance monitoring can be less intensive (such as with longer intervals between performance checks), and/or that the retraining is no longer necessary (at least for a period of time, such as until a long-term maintenance period arrives and/or until there are significant shifts in context of utilization). As the ANN continues to improve upon existing functions and/or add new functions via the dual-process training process, the ANN may perform other, at times more "intellectually-demanding" (e.g.,

retraining intensive) tasks simultaneously. For example, utilizing dual process-learned knowledge of a function or process being trained, the ANN can solve an unrelated complex problem or make a retraining decision simultaneously. The retraining may include supervision, such as where an agent (e.g., human supervisor or intelligent agent) directs the ANN to a retraining objective (e.g., "master this new function") and provides a set of training tasks and feedback functions (such as supervisory grading) for the retraining. In-embodiments, the ANN can be used to organize the supervision, training and retraining of other dual process-trained ANNs, to seed such training or retraining, or the like.

[2679] In example embodiments, one or more behaviors and operational processes (such as decision-making) of the ANN may be products of training and retraining processes facilitated by the training system and the retraining system, respectively. The training system may be configured to perform automatic training of ANN, such as by continuously adding additional instances of training data as it is collected by or from various data sources. The retraining system may be configured to perform effortful, analytical, intentional retraining of the ANN, such as based on memory (e.g., stored training data or refined training data) and/or optionally based on reasoning or other factors. For example, in a deployment management context, the training system may be associated with a standard response by the ANN, while the retraining system may implement DPLF retraining and/or network adaptation of the ANN. In some cases, retraining of the ANN beyond the factory, or "out-of-the-box," training level may involve more than retraining by the retraining system. Successful adjustment of the ANN by one or more network adaptations may be dependent on the operation of one or more network adjustments of the training system.

[2680] In example embodiments, the training system may facilitate fast operating by and training of the ANN by applying existing neural functions of the ANN based on training of the ANN with previous datasets. Standard operational activities of the ANN that may draw heavily on the training system may include one or more of the methods, processes, workflows, systems, or the like described throughout this disclosure and the documents incorporated herein, such as, without limitation: defined functions within networking (such as discovering available networks and connections, establishing connections in networks, provisioning network bandwidth among devices and systems, routing data within networks, steering traffic to available network paths, load balancing across networking resources, and many others); recognition and classification (such as of images, text, symbols, objects, video content, music and other audio content, speech content, and many others); spoken words; prediction of states and events (such as prediction of failure modes of machines or systems, prediction of events within workflows, predictions of behavior in shopping and other activities, and many others); control (such as controlling autonomous or semi-autonomous systems, automated agents (such as automated call-center operations, chat bots, and the like) and others); and/or optimization and recommendation (such as for products, content, decisions, and many others). ANNs may also be suitable for training datasets for scenarios that only require output. The standard operational activities may not require the ANN to actively analyze what is being asked of the ANN beyond operating on well-defined data inputs, to calculate well-defined outputs for well-defined use cases. The operations of the training system and/or the

retraining system may be based on one or more historic data training datasets and may use the parameters of the historic data training datasets to calculate results based on new input values and may be performed with small or no alterations to the ANN or its input types. In example embodiments, an instance of the training system can be trained to classify whether the ANN is capable of performing well in a given situation, such as by recognizing whether an image or sound being classified by the ANN is of a type that has historically been classified with a high accuracy (e.g., above a threshold).

[2681] In example embodiments, network adaptation of the ANN by one or both of the training system and the retraining system may include a number of defined network functions, knowledge, and intuition-like behavior of the ANN when subjected to new input values. In such embodiments, the retraining system may apply the new input values to the DPLF system to adjust the functional response of the ANN, thereby performing retraining of the ANN. The DPANN system 70000 may determine that retraining the ANN via network adjustment is necessary when, for example, without limitation, functional neural networks are assigned activities and assignments that require the ANN to provide a solution to a novel problem, engage in network adaptation or other higher-order cognitive activity, apply a concept outside of the domain in which the DPANN was originally designed, support a different context of deployment (such as where the use case, performance requirements, available resources, or other factors have changed), or the like. The ANN can be trained to recognize where the retraining system is needed, such as by training the ANN to recognize poor performance of the training system, high variability of input data sets relative to the historical data sets used to train the training system, novel functional or performance requirements, dynamic changes in the use case or context, or other factors. The ANN may apply reasoning to assess performance and provide feedback to the retraining system. The ANN may be trained and/or retrained to perform intuitive functions, optionally including by a combinatorial or re-combinatorial process (e.g., including genetic programming wherein inputs (e.g., data sources), processes/functions (e.g., neural network types and structures), feedback, and outputs, or elements thereof, are arranged in various permutations and combinations and the ANN is tested in association with each (whether in simulations or live deployments), such as in a series of rounds, or evolutionary steps, to promote favorable variants until a preferred ANN, or preferred set of ANNs is identified for a given scenario, use case, or set of requirements). This may include generating a set of input "ideas" (e.g., combinations of different conclusions about cause-and-effect in a diagnostic process) for processing by the retraining system and subsequent training and/or by an explicit reasoning process, such as a Bayesian reasoning process, a casuistic or conditional reasoning process, a deductive reasoning process, an inductive reasoning process, or others (including combinations of the above) as described in this disclosure or the documents incorporated herein by reference.

[2682] Referring to Figure 258 in example embodiments, the DPLF may perform an encoding process of the DPLF to process datasets into a stored, form for future use, such as retraining of the ANN by the retraining system. The encoding process enables datasets to be taken in, understood, and altered by the DPLF to better support storage in and usage from the memory. The DPLF may

apply current functional knowledge and/or reasoning to consolidate new input values. The memory can include short-term memory (STM), long-term memory (LTM), or a combination thereof. The datasets may be stored in one or both of the STM and the LTM. The STM may be implemented by the application of specialized behaviors inside the ANN (such as recurrent neural network, which may be gated or un-gated, or long-term short-term neural networks). The LTM may be implemented by storing scenarios, associated data, and/or unprocessed data that can be applied to the discovery of new scenarios. The encoding process may include processing and/or storing, for example, visual encoding data (e.g., processed through a Convolution Neural Network), acoustic sensor encoding data (e.g., how something sounds, speech encoding data (e.g., processed through a deep neural network (DNN), optionally including for phoneme recognition), semantic encoding data of words, such to determine semantic meaning, e.g., by using a Hidden Markov Model (HMM); and/or movement and/or tactile encoding data (such as operation on vibration/accelerometer sensor data, touch sensor data, positional or geolocation data, and the like). While datasets may enter the DPLF system through one of these modes, the form in which the datasets are stored may differ from an original form of the datasets and may pass-through neural processing engines to be encoded into compressed and/or context-relevant format. For example, an unsupervised instance of the ANN can be used to learn the historic data into a compressed format.

[2683] In example embodiments, the encoded datasets are retained within the DPLF system. Encoded datasets are first stored in short-term DPLF, i.e., STM. For example, sensor datasets may be primarily stored in STM, and may be kept in STM through constant repetition. The datasets stored in the STM are active and function as a kind of immediate response to new input values. The DPANN system 70000 may remove datasets from STM in response to changes in data streams due to, for example, running out of space in STM as new data is imported, processed and/or stored. For example, it is viable for short-term DPLF to only last between 15 and 30 seconds. STM may only store small amounts of data typically embedded inside the ANN.

[2684] In example embodiments, the DPANN system 70000 may measure attention based on utilization of the training system, of the DPANN system 70000 as a whole, and/or the like, such as by consuming various indicators of attention to and/or utilization of outputs from the ANN and transmitting such indicators to the ANN in response (similar to a "moment of recognition" in the brain where attention passes over something and the cognitive system says "aha!"). In example embodiments, attention can be measured by the sheer amount of the activity of one or both of the systems on the data stream. In example embodiments, a system using output from the ANN can explicitly indicate attention, such as by an operator directing the ANN to pay attention to a particular activity (e.g., to respond to a diagnosed problem, among many other possibilities). The DPANN system 70000 may manage data inputs to facilitate measures of attention, such as by prompting and/or calculating greater attention to data that has high inherent variability from historical patterns (e.g., in rates of change, departure from norm, etc.), data indicative of high variability in historical performance (such as data having similar characteristics to data sets involved in situations where the ANN performed poorly in training), or the like.

[2685] In example embodiments, the DPANN system 70000 may retain encoded datasets within the DPLF system according to and/or as part of one or more storage processes. The DPLF system may store the encoded datasets in LTM as necessary after the encoded datasets have been stored in STM and determined to be no longer necessary and/or low priority for a current operation of the ANN, training process, retraining process, etc. The LTM may be implemented by storing scenarios, and the DPANN system 70000 may apply associated data and/or unprocessed data to the discovery of new scenarios. For example, data from certain processed data streams, such as semantically encoded datasets, may be primarily stored in LTM. The LTM may also store image (and sensor) datasets in encoded form, among many other examples.

[2686] In example embodiments, the LTM may have relatively high storage capacity, and datasets stored within LTM may, in some scenarios, be effectively stored indefinitely. The DPANN system 70000 may be configured to remove datasets from the LTM, such as by passing LTM data through a series of memory structures that have increasingly long retrieval periods or increasingly high threshold requirements to trigger utilization (similar to where a biological brain "thinks very hard" to find precedent to deal with a challenging problem), thereby providing increased salience of more recent or more frequently used memories while retaining the ability to retrieve (with more time/effort) older memories when the situation justifies more comprehensive memory utilization. As such, the DPANN system 70000 may arrange datasets stored in the LTM on a timeline, such as by storing the older memories (measured by time of origination and/or latest time of utilization) on a separate and/or slower system, by penalizing older memories by imposing artificial delays in retrieval thereof, and/or by imposing threshold requirements before utilization (such as indicators of high demand for improved results). Additionally or alternatively, LTM may be clustered according to other categorization protocols, such as by topic. For example, all memories proximal in time to a periodically recognized person may be clustered for retrieval together, and/or all memories that were related to a scenario may be clustered for retrieval together.

[2687] In example embodiments, the DPANN system 70000 may modularize and link LTM datasets, such as in a catalog, a hierarchy, a cluster, a knowledge graph (directed/acyclic or having conditional logic), or the like, such as to facilitate search for relevant memories. For example, all memory modules that have instances involving a person, a topic, an item, a process, a linkage of n-tuples of such things (e.g., all memory modules that involve a selected pair of entities), etc. The DPANN system 70000 may select sub-graphs of the knowledge graph for the DPLF to implement in one or more domain-specific and/or task-specific uses, such as training a model to predict robotic or human agent behavior by using memories that relate to a particular set of robotic or human agents, and/or similar robotic or human agents. The DPLF system may cache frequently used modules for different speed and/or probability of utilization. High value modules (e.g., ones with high-quality outcomes, performance characteristics, or the like) can be used for other functions, such as selection/training of STM keep/forget processes.

[2688] In example embodiments, the DPANN system 70000 may modularize and link LTM datasets, such as in various ways noted above, to facilitate search for relevant memories. For example, memory modules that have instances involving a person, a topic, an item, a process, a

linkage of n-tuples of such things (such as all memory modules that involve a selected pair of entities), or all memories associated with a scenario, etc., may be linked and searched. The DPANN system 70000 may select subsets of the scenario (e.g., sub-graphs of a knowledge graph) for the DPLF for a domain-specific and/or task-specific use, such as training a model to predict robotic or human agent behavior by using memories that relate to a particular set of robotic or human agents and/or similar robotic or human agents. Frequently used modules or scenarios can be cached for different speed/probability of utilization, or other performance characteristics. High value modules or scenarios (ones where high-quality outcomes results) can be used for other functions, such as selection/training of STM keep/forget processes, among others.

[2689] In example embodiments, the DPANN system 70000 may perform LTM planning, such as to find a procedural course of action for a declaratively described system to reach its goals while optimizing overall performance measures. The DPANN system 70000 may perform LTM planning when, for example, a problem can be described in a declarative way, the DPANN system 70000 has domain knowledge that should not be ignored, there is a structure to a problem that makes the problem difficult for pure learning techniques, and/or the ANN needs to be trained and/or retrained to be able to explain a particular course of action taken by the DPANN system 70000. In example embodiments, the DPANN system 70000 may be applied to a plan recognition problem, i.e., the inverse of a planning problem: instead of a goal state, one is given a set of possible goals, and the objective in plan recognition is to find out which goal was being achieved and how.

[2690] In example embodiments, the DPANN system 70000 may facilitate LTM scenario planning by users to develop long-term plans. For example, LTM scenario planning for risk management use cases may place added emphasis on identifying extreme or unusual, yet possible, risks and opportunities that are not usually considered in daily operations, such as ones that are outside a bell curve or normal distribution, but that in fact occur with greater-than-anticipated frequency in "long tail" or "fat tail" situations, such as involving information or market pricing processes, among many others. LTM scenario planning may involve analyzing relationships between forces (such as social, technical, economic, environmental, and/or political trends) in order to explain the current situation, and/or may include providing scenarios for potential future states.

[2691] In example embodiments, the DPANN system 70000 may facilitate LTM scenario planning for predicting and anticipating possible alternative futures along with the ability to respond to the predicted states. The LTM planning may be induced from expert domain knowledge or projected from current scenarios, because many scenarios (such as ones involving results of combinatorial processes that result in new entities or behaviors) have never yet occurred and thus cannot be projected by probabilistic means that rely entirely on historical distributions. The DPANN system 70000 may prepare the application to LTM to generate many different scenarios, exploring a variety of possible futures to the DPLF for both expected and surprising futures. This may be facilitated or augmented by genetic programming and reasoning techniques as noted above, among others.

[2692] In example embodiments, the DPANN system 70000 may implement LTM scenario planning to facilitate transforming risk management into a plan recognition problem and apply the

DPLF to generate potential solutions. LTM scenario induction addresses several challenges inherent to forecast planning. LTM scenario induction may be applicable when, for example, models that are used for forecasting have inconsistent, missing, unreliable observations; when it is possible to generate not just one but many future plans; and/or when LTM domain knowledge can be captured and encoded to improve forecasting (e.g., where domain experts tend to outperform available computational models). LTM scenarios can be focused on applying LTM scenario planning for risk management. LTM scenarios planning may provide situational awareness of relevant risk drivers by detecting emerging storylines. In addition, LTM scenario planning can generate future scenarios that allow DPLF, or operators, to reason about, and plan for, contingencies and opportunities in the future.

[2693] In example embodiments, the DPANN system 70000 may be configured to perform a retrieval process via the DPLF to access stored datasets of the ANN. The retrieval process may determine how well the ANN performs with regard to assignments designed to test recall. For example, the ANN may be trained to perform a controlled vehicle parking operation, whereby the autonomous vehicle returns to a designated spot, or the exit, by associating a prior visit via retrieval of data stored in the LTM. The datasets stored in the STM and the LTM may be retrieved by differing processes. The datasets stored in the STM may be retrieved in response to specific input and/or by order in which the datasets are stored, e.g., by a sequential list of numbers. The datasets stored in the LTM may be retrieved through association and/or matching of events to historic activities, e.g., through complex associations and indexing of large datasets.

[2694] In example embodiments, the DPANN system 70000 may implement scenario monitoring as at least a part of the retrieval process. A scenario may provide context for contextual decision-making processes. In example embodiments, scenarios may involve explicit reasoning (such as cause-and-effect reasoning, Bayesian, casuistic, conditional logic, or the like, or combinations thereof) the output of which declares what LTM-stored data is retrieved (e.g., a timeline of events being evaluated and other timelines involving events that potentially follow a similar cause-and-effect pattern). For example, diagnosis of a failure of a machine or workflow may retrieve historical sensor data as well as LTM data on various failure modes of that type of machine or workflow (and/or a similar process involving a diagnosis of a problem state or condition, recognition of an event or behavior, a failure mode (e.g., a financial failure, contract breach, or the like), or many others).

[2695] FIG. 259 is a diagrammatic view illustrating an example implementation of a conventional computer vision system 71100 for recognizing an object 71102 of interest. The computer vision system 71100 includes a lens assembly 71104 that attempts to focus light from the object 71102 onto a sensor 71106. The sensor 71106 may be an image sensor such as a charge coupled device (CCD) or complementary metal oxide semiconductor (CMOS) device containing array of photo sensitive elements. The sensor may convert the light into analog electrical signal corresponding to light intensity. An analog to digital (AD) converter 71108 then converts analog voltage into digital data. This raw digital data is then sent to an image processing system 71110 for analysis. The image processing system 71110 then processes the raw digital data to generate an image 71112. The

image processing system 71110 may also involve pre-processing and post-processing including image scaling, noise reduction, color adjustment, brightness adjustment, white balance adjustment, sharpness, adjustment, contrast adjustment and the like to enhance the image quality. Further the image may be analyzed using machine learning or other algorithms to identify one or more objects in the image.

[2696] Conventional computer vision systems 71100 have many limitations. The attempt to recreate vision by creating focused images leads to the loss of a large amount of information and leaves the vision system 71100 with limited data. The computer vision system 71100 typically generates two dimensional images of three-dimensional objects and is unable to capture information related to aspects like object depth, motion, orientation and the like. The algorithms in the computer vision system 71100 attempt to infer information about a 3D scene/object from 2D frames and information thereby limiting the quality of inferences.

[2697] FIG. 260 is a schematic illustrating an example implementation of a dynamic vision system 71200 for dynamically learning an object concept about an object 71202 of interest according to an embodiment of the present disclosure. The dynamic vision system 71200 may replace and/or augment the lens 71104 of a conventional vision system 71100 with a variable focus liquid lens 71204. The variable focus liquid lens 71204 may be an electrically controlled cell containing optical-grade liquid, that is deformed through electric current, changing the shape of the lens. The dynamic vision system 71200 leverages this flexibility of liquid lens 71204 by constantly adjusting lens parameters to dynamically change various optical characteristics of light that pass through the lens including focal length, spherical aberration, field curvature, coma, chromatics aberrations, distortion, vignetting, ghosting and flaring, and diffraction of light. A fully variable liquid lens thus allows for more dynamic input for a sensor 71206 enabling it to capture visual information and metadata that is otherwise lost in the conventional computer vision system 71100.

[2698] An analog to digital (AD) converter 71208 may generate digital data from the rich visual information captured at the sensor 71206 and an image processing system 71208 with pre-processing, and post-processing capabilities may generate images that are based with additional optical parameters as part of image. The processing system 71209 may also include a control system 71212 configured to adjust one or more optical parameters in real time including focal length, liquid materials, specularity, color, environment and lens shape. An adaptive intelligence system 71214 may then dynamically learn on a training set of outcomes, parameters, and data collected from the liquid lens 71204 to generate an object concept 71216. The object concept 71216 may include contextual intelligence about the object and its environment which may then be processed by adaptive intelligence system 71214 to recognize the object 71202.

[2699] In example embodiments, the adaptive intelligence system 71214 may include artificial intelligence capability, such as involving machine learning or other algorithms, neural networks, expert systems, models and others, to process the input data from the liquid lens and dynamically learn the object concept to provide superior object recognition and vision.

[2700] In example embodiments, the dynamic vision system 71200 may feed real-time adjustable data streams to the processing system 71209 to generate situational awareness or create out-of-

focus images of the object 71202 so as to capture large amounts of information that is otherwise lost when inferring depth and distance in a focused image of a conventional vision system 71100. The dynamic input to the liquid lens 71204 may provide richer metadata for image processing as the images are based on additional optical parameters than just focal length and aperture. The image processing system 71210 may incorporate previously lost information so as to generate a new set of insights about the object and its surroundings not captured by the conventional computer vision systems 71100.

[2701] Compared to conventional computer vision systems 71100, that utilize fixed sensory elements, the dynamic vision system 71200 provided herein may utilize a dynamically learned liquid lens assembly. The conformable liquid lens 71204 in the assembly may be continuously, and/or frequently, adjusting based on, for example, environment factors and/or on feedback from the processing system 71209 to generate training data that is deeper in context and that corresponds to the physical light that the image represents. By training the dynamic vision system 71200 to recognize objects using variable optical parameters through the liquid lens assembly, the processing system 71209 may learn an optimum optical setting(s) to detect an object. The more dynamic input to the dynamic vision system 71200 may result in creating a richer context and providing superior object recognition.

[2702] The dynamic vision system 71200 may integrate sensing, control and processing functions and dynamically adjusts the liquid lens 71200 as the vision algorithms in the processing system 71209 take different inputs to produce a real-world vision result.

[2703] The dynamic vision system 71200 mimics biological vision by integrating sensing, control and processing functions (biological vision involves a stream of information passing directly through deep learning systems where these deep learning systems can directly change aspects of vision processing, including orientation, fovea centralis attention, eyelid actions, blinking and communication with other humans).

[2704] In example embodiments, the dynamic vision system 71200 may utilize saccades to characterize objects by context and build a rich model of the object in its environment by capturing contextual intelligence through associations. This mirrors how saccades capture information about an object in its environment. Saccade denotes a quick, simultaneous movement of both eyes between two or more areas of focus. While viewing a scene, human eyes make sporadic saccadic movements stopping several times while locating key parts of the scene, moving quickly between each stop and building up a mental three-dimensional map corresponding to the scene. The dynamic vision system 71200 and methods described herein may use saccades to characterize objects by context and allow control of an optical system to more quickly identify and characterize a field of view. Saccades integrate varying physical/optical properties, along with object-oriented learning, to rapidly improve understanding and search in the visual sphere.

[2705] In example embodiments, the dynamic vision system 71200 may also mimic biofeedback loops of human babies to create a system of associative memory and vision and build a causal three-dimensional model of the environment. The learning system in human babies involves many feedback loops of activities wherein babies build a causal model of the world around them by

performing sequences of controlled experiments. The dynamic vision offered by the liquid lens-based vision system may, in part, mirror the learning algorithm of babies by starting a training set around the object and letting its learning algorithm figure out the right way to look at the object.

[2706] FIG. 261 depicts a schematic illustrating an example architecture of a dynamic vision system 71300 depicting a detailed view of various components according to some embodiments of the present disclosure. The dynamic vision system 71300 for recognizing an object 71302 may include an optical assembly 71304 and a processing system 71306. The optical assembly 71304 may include a conformable liquid lens 71308, one or more sensors 71310 and an analog to digital (AD) converter 71312. The processing system 71306 may include a control system 71314, an image processing system 71316, an adaptive intelligence system 71318, a digital twin system 71320 and a simulation system 71322. The adaptive intelligence system may include a machine learning system 71324 and an artificial intelligence system 71326.

[2707] The conformable liquid lens 71308 of the optical assembly 71304 may frequently adjust in real-time based, in part, on change in one or more optical parameters by the control system 71314 creating real-time data streams at the sensors 71310 which are then provided to the processing system 71306 to generate a situational awareness or computerized understanding of the world that the dynamic vision system 71300 is operating in. This understanding may include rich contextual intelligence about the object and its environment and may be represented as an object concept. The object concept may be used by the processing system for object recognition, predicting object motion, location and orientation, creating a 3D model of the object, monitoring the object for any defects and other applications. For example, the adaptive intelligence system 71318 may process the object concept to build a three-dimensional representation of the object. The machine learning system 71322 in the adaptive intelligence system 71318 may input the object concept into one or more machine learning models, the object concept being used as training data for the machine learning models. Further, the artificial intelligence system 71324 may be configured to make classifications, predictions, and other decisions relating to the object including determining the position, orientation and motion of the object.

[2708] In example embodiments, the dynamic vision system 71300 may be configured to process sensor information to create a three-dimensional representation of the object 71302 in a single step without the intermediate step of processing into flat images.

[2709] In example embodiments, the control system 71314 may provide control instructions to one or more actuators which in turn drive the adjustments in liquid lens configurations. The actuators may be operated by a source of energy, typically electric current, hydraulic fluid pressure, or pneumatic pressure, and convert that energy into motion. Examples of actuators may include linear actuators, solenoids, comb drives, digital micromirror devices, electric motors, electroactive polymers, hydraulic cylinders, piezoelectric actuators, pneumatic actuators, servomechanisms, servo motors, thermal bimorphs, screw jacks, or any other type of hydraulic, pneumatic, electric, mechanical, thermal, magnetic type of actuator, or some other type of actuator.

[2710] In example embodiments, the control system 71314 may provide control instructions to one or more actuators to change the focal length of the liquid lens based on stimulation. This may

provide the dynamic vision system 71300 with an auto-focus capability by focusing, refocusing or defocusing the lens to a desired focal length. The stimulation mechanism may include electrical, hydraulic, pneumatic, mechanical, thermal or magnetic.

[2711] Some examples of control systems 71314 include electrowetting, sound piezoelectrics and electro-active polymers.

[2712] In example embodiments, the conformable liquid lens assembly in the dynamic vision system 71300 may have an electrowetting control system such that an application of electrical voltage to the fluid in the liquid lens changes the shape of the liquid, effectively changing the focus of the liquid lens assembly.

[2713] In example embodiments, the placement of actuators in a variable focused liquid lens based optical assembly may be optimized using machine learning.

[2714] In example embodiments, the control system 71314 may control the liquid lens 71304 configuration based on feedback from the processing system 71306 in response to a change in environment factors. Some examples of the environmental factors include temperature, vibrations, ambient sensor data, workflows, entity IDs, user behavioral data, entity profiling, similarity to known data and the like.

[2715] In example embodiments, the control system 71314 may control the liquid lens 71304 configuration based on feedback from the processing system 71306 in response to a change in source lighting including control color, color temperature, timing (PWM), amplitude (e.g., increase PWM but diminish amplitude, direction, polarization, and the like.

[2716] In example embodiments, the control system 71314 may control the liquid lens configuration based on human occupancy and awareness of when lighting needs to be coordinated with human needs versus adjusted solely to serve the liquid lens system.

[2717] In example embodiments, the optical assembly 71304, may include multiple sets of liquid lenses with the processing system 71306 coordinating the control of multiple liquid lenses setup.

[2718] In example embodiments, the optical assembly 71304, may include multiple sets of liquid lenses with each lens having a separate objective function, and a separate processing system with AI setups or algorithms.

[2719] In example embodiments, the optical assembly 71304, may include one or more liquid lens combined with a conventional convex or concave optical lens with the processing system 71306 coordinating the control of the combination.

[2720] In example embodiments, the processing system 71306, such as using adaptive intelligence system 71318, the digital twin system 71320 and the simulation system 71322 may execute simulations to model, simulate and characterize the mechanical, optical, or lighting aspects of the dynamic vision system 71300. The simulations executed by the processing system 71306 may help identify suitable imaging components for the dynamic vision system 71300 including sensors, lenses and lights. The simulations may include real time analytics to calculate a wide range of metrics, build charts, graphs and models and visualize the effect of change of one or more optical parameters on the performance of the dynamic vision system 71300. The artificial intelligence system 71326 in the adaptive intelligence system 71318 may then utilize the one or more models

to make classifications, predictions, recommendations, and/or to generate or facilitate decisions or instructions relating to the lens materials, geometry, optical properties, performance and design of the dynamic vision system 71300. For example, the artificial intelligence system 71326 may execute simulations on one or more liquid lens digital twins for generating recommendations relating to the fluid used in the liquid lens. The simulations may be performed using different fluids including distilled water, methyl alcohol, ethyl alcohol, ether, carbon tetrachloride, methyl acetate, glycerine, nitrobenzene and the like to generate recommendations on the preferred fluid for a given application of the dynamic vision system 71300.

[2721] The dynamic vision system 71300 may utilize dynamically learned sensory elements to recognize objects ensuring a richer object recognition capacity that may be applied to a very wide range of use cases. The approach is ideal for imaging applications requiring rapid focusing, high throughput, and depth of field and working distance accommodation. Moreover, the approach is especially beneficial for complex vision applications where conventional vision technologies have been inadequate. Some examples of such applications include: recognizing objects in dynamic environments like when the object or vision system are moving; recognizing three dimensional (3D) objects by capturing depth data; recognizing tiny objects; recognizing facial features; recognizing novel or previously unseen objects, recognizing objects in a power constrained or network constrained environment; and so on.

[2722] In example embodiments, the dynamic vision system 71300 may utilize associative learning to recognize novel or previously unseen objects (i.e., objects that were not part of the training data set). Associative learning recognizes objects by accessing an attribute layer that includes attribute recognition programs (e.g., one for recognizing dimensions, another for shape, a third for color, a fourth for material, and so on). An attribute may be an inherent characteristic of an object and may be defined in terms of appearance adjectives, such as dimensions, color, shape, pattern, texture, material, and the like, the presence and absence of parts (e.g., has legs, has wheels, has a head), and similarity to known objects (e.g., similar to chair, sofa, table, car). Thus, the dynamic vision system 71300 may learn about an object in multiple streams of attributes constituting the object by training classifiers to recognize each object attribute individually. For example, a cup may be learned in multiple streams of attributes and artificial intelligence classifiers may be trained to recognize each attribute individually. The first stream may relate to material composition of the cup (made of glass, wood, metal, ceramic), a second stream may relate to color (red, blue), a third stream may relate to presence and absence of parts (has a handle), a fourth stream may relate to similarity to known objects (similar to containers, boxes, automobiles). Once the attributes have been learned individually, the dynamic vision system 71300 may search against a set of stored objects that have attributes similar to the cup for recognition.

[2723] In example embodiments, the dynamic vision system 71300 may be integrated into or with a set of robotic systems, such as mobile and/or autonomous robotic systems. For example, the dynamic vision system 71300 may be contained within the housing or body of a robotic system, such as a multi-purpose/general purpose robotic system, such as one that simulates human or other animal species capabilities. The vision capabilities may enable the robot in identifying and

manipulating a target object for use in robotic assembly lines where object depth, orientation, position and motion may be inferred for improved object identification. The vision capabilities may also enable the robot in simultaneous localization and mapping, which is a technique for estimating the position of the robot with respect to its surroundings while mapping the environment at the same time. As another example, the dynamic vision system 71300 may be integrated with a robotic exoskeleton designed to augment the capabilities of a human operator and provide optimized sensing and control for the human operator.

[2724] In example embodiments, the output from the dynamic vision system 71300 may be temporally combined with output from other sensors in the robot using conditional probabilities to create a combined view of the object that is richer and includes information about the position, orientation and motion of the object. Some examples of sensors that may be used in conjunction with the liquid lens based dynamic vision system 71300 include cameras, LIDARs, RADARs, SONARs, thermal imaging sensor, hyperspectral imaging sensor, illuminance sensors, force sensors, torque sensors, velocity sensors, acceleration sensors, position sensors, proximity sensors, gyro sensors, sound sensors, motion sensors, location sensors, load sensors, temperature sensors, touch sensors, depth sensors, ultrasonic range sensors, infrared sensors, chemical sensors, magnetic sensors, inertial sensors, gas sensors, humidity sensors, pressure sensors, viscosity sensors, flow sensors, object sensors, tactile sensors, or some other type of sensor.

[2725] In example embodiments, the dynamic vision system 71300 incorporating a conformable liquid lens controlled by AI as necessary, and augmented by sensors may be adapted to build a neural prosthetics system.

[2726] In example embodiments, the dynamic vision system 71300 incorporating a conformable liquid lens technology controlled by AI as necessary, may be adapted to build an exoskeleton system.

[2727] In example embodiments, the dynamic vision system 71300 incorporating a conformable liquid lens controlled by AI as necessary, and augmented by sensors may be adapted to perform facial recognition for human faces obscured by face masks.

[2728] FIG. 262 depicts a flow diagram illustrating a method for object recognition by the liquid lens based dynamic vision system according to some embodiments of the present disclosure.

[2729] Referring to Fig 4, at 71402, real time data streams representing object concept are received from the liquid lens based optical assembly. The data streams may be received at the sensor and include rich contextual and visual information generated by constantly adjusting liquid lens in response to changes in optical parameters. The data streams may be analyzed at edge devices or sent to data processing by local or remote intelligence. The use of cloud-connectable edge devices, such as within computing infrastructure that is proximal to the dynamic vision system 71300 and/or that is integrated with or into the dynamic vision system 71300, such as where the dynamic vision system 71300 has onboard edge computational and/or connectivity resources, such as 5G (or other cellular), Wi-Fi, BLUETOOTH™, fixed networking resources, or the like, may offer opportunities to provide rapid, real-time or near real-time processing responsiveness. At 71404, the real-time data streams are processed by the image processing system to determine an object concept that

includes contextual intelligence about the object and its environment. At 71406, the optical parameters are adjusted by the control system leading to a change in configuration of the liquid lens. The constantly adjusting liquid lens creates real time data streams at the sensor and rich metadata for image processing as the images are based on additional optical parameters than just focal length and aperture. At 71408, the object concept is sequentially revised and used as an input to train a machine learning model, which dynamically learns on a training set of outcomes, parameters and data collected from the liquid lens based optical assembly. At 71410, the object concept including contextual intelligence about the object and its environment is utilized by artificial intelligence to make classifications, predictions, and other decisions relating to the object including determining the position, orientation and motion of the object.

[2730] FIG. 263 depicts a schematic illustrating an example implementation of a dynamic vision system for modelling, simulating and optimizing various optical, mechanical, design and lighting parameters of the dynamic vision system according to some embodiments of the present disclosure. The dynamic vision system may learn on data captured by sensors in response to sequentially adjusting the liquid lens to train the artificial learning system to use digital twins for classification, predictions and decision-making.

[2731] The digital twin system 71320 may be configured to simulate operation of the dynamic vision system 71300 so as to continuously capture the key operational metrics and may be used to monitor and optimize the performance of the dynamic vision system 71300 in real-time, or near real-time. The digital twin system 71320 may create a digital replica or digital twins 71502 of one or more of the components or subsystems of the dynamic vision system 71300. The digital twins 71502 of the one or more of the components or subsystems may use substantially real-time sensor data to provide for substantially real-time virtual representation and for simulation of one or more possible future states of the one or more components and subsystems. The digital twins 71502 may be updated continuously based on sensor data, to reflect the current condition or parameter values of the component or subsystem. The digital twins thus provide a high fidelity, digital simulation of the behavior of the component or subsystem. This capability may be used to produce a digital profile of both the prior and current behaviors of the component or subsystem with the resulting profile used to detect behavior that is less than optimal as well as to predict future behavior of the component or subsystem.

[2732] In example embodiments, attribute twins may be defined for one or more objects to be recognized by the dynamic vision system 71300. The attribute twins may be "pseudo digital twins" spanning across multiple objects and focused on learning the fundamental concepts of object detection. For example, the attribute twin "shape" may denote the shape of an object and learn about shape in all its forms and possibilities. Similarly, the attribute twin "material" may learn about various material compositions of various objects and a "color" digital twin learn about various colors. To extend the example of recognizing a cup described above, the attribute twins may help the dynamic vision system 71300 recognize a novel blue metal cup (not in the training data set) by identifying that the novel object or cup is blue in "color", made of "metal" and is shaped as a "cylindrical container". This can also lead to new pseudo digital twins that not only

allow for discovering completely novel objects, but also allows the dynamic vision system 71300 to then determine a relevant "pseudo digital action" associated with the recognized object. For example, some pseudo digital actions associated with the cup may be "lift", "drop", "grasp", "serve" and the like.

[2733] Referring to FIG. 5, the digital twins 71502 in the dynamic vision system 71300 may include object twin 71504, environment twin 71506, liquid lens twin 71508, optical lens twin 71510, sensor twin 71512, process twin 71514, actuator twin 71516, object concept twin 71518 and the like, that allow for modeling, simulation, prediction, decision-making, and classification by the processing system 71306. The digital twins 71502 may be populated with relevant data, for example the liquid lens twins 71508 may be populated with data related to corresponding a liquid lens including dimension data, material data, shape data, feature data, thermal data, vibration data, and the like. The digital twins may provide one or more simulations of both physical elements and characteristics of the one or more components or subsystems being replicated and the dynamics thereof, in embodiments throughout the lifecycle of the one or more components being replicated.

[2734] In example embodiments, the digital twins 71502 may provide a hypothetical simulation of the one or more components or subsystems, for example during a design phase before the one or more components are manufactured or fabricated, or during or after construction or fabrication of the one or more components by allowing for hypothetical extrapolation of sensor data to simulate a state of the one or more components, such as during any suitable hypothetical situation. In example embodiments, the machine learning model 71520 may automatically predict hypothetical situations for simulation with the digital twins 71502, such as by predicting possible improvements to the one or more components, predicting if one or more components are compatible with one another, predicting when one or more components may fail and/or suggesting possible improvements to the one or more components, such as changes to parameters, arrangements, configurations, or any other suitable change to the components. For example, the liquid lens twins 71506 and optical lens twins 71510 may run hypothetical simulations to check for compatibility with one another as well as with the optical assembly and predict the optimal arrangement in the assembly.

[2735] In example embodiments, the machine learning models 71520 in conjunction with digital twins 71502 may help drive various applications including material selection 71522, design optimization 71524, and motion prediction 71526.

[2736] In example embodiments, the digital twins 71502 may allow for simulation of the one or more components during both design and operation phases of the one or more components, as well as simulation of hypothetical operation conditions and configurations of the one or more components by facilitating observation, measurement and analysis of various metrics and then passing the insights onto the design or operational processes for improvement of these processes.

[2737] The simulation system 71322 may set up, provision, configure, and otherwise manage interactions and simulations between and among digital twins 71502. Thus, the simulation system may help simulate, evaluate and optimize the behavior and characteristics of various components

and subsystems of the dynamic vision system 71300 using the digital twins 71502 of such components and subsystems.

[2738] In example embodiments, the artificial intelligent system 71326 may be configured to execute simulations in the simulation system 71322 using the liquid lens twins 71508 and/or other digital twins 71502 available to the digital twin system 71214. For example, the processing system 71306 may adjust one or more optical parameters of the liquid lens twin 71508. In example embodiments, the artificial intelligent system 71326 may, for each set of parameters, execute a simulation based on the set of parameters and may collect the simulation outcome data resulting from the simulation. For example, the artificial intelligent system 71326 may execute simulations by varying the optical parameters of the liquid lens twin 71506 to generate simulation outcomes in the form of object concept twin 71518. During the simulation, the processing system 71306 may vary the focal length, fluid materials, specularity, color, environment, lens shape and any other parameters of the liquid lens twin 71506. The outcome data from such simulations in the form of object concept twins 71518 in addition to other sensor data as well as data from other sources may then be used to train the machine learning models 71520 by the machine learning system 71324.

[2739] In example embodiments, the machine learning models 71520 may process the data received from sensors, including the event data and the state data to define simulation data for use by the digital twin system 71320. The machine learning models 71520 may, for example, receive state data and event data related to a particular component of the dynamic vision system 71300 and perform a series of operations on the state data and the event data to format the state data and the event data into a format suitable for use by the digital twin system 71320. For example, machine learning models 71520 may collect data from one or more sensors positioned on, near, in, and/or around the liquid lens to process the sensor data into simulation data and output the simulation data to the digital twin system 71320. The digital twin system 71320 may then use the simulation data to create the liquid lens twin 71506, the simulation including for example metrics including shape, material, focal length, specularity, environment, lighting, color, temperature, pressure, wear and vibration. The simulation may be a substantially real-time simulation, allowing for a user of the dynamic vision system 71300 to view the simulation of the liquid lens, metrics related thereto, and metrics related to parts thereof, in substantially real time. The simulation may be a predictive or hypothetical situation, allowing for a user of the dynamic vision system 71300 to view a predictive or hypothetical simulation of the liquid lens, metrics related thereto, and metrics related to components thereof.

[2740] In example embodiments, the machine learning models 71520 and the digital twin system 71320 may process sensor data and create a digital twin for a set of components to facilitate real-time simulation, predictive simulation, and/or hypothetical simulation of a related group of components.

[2741] The machine learning models 71520 may be algorithms and/or statistical models that perform specific tasks without using explicit instructions, relying instead on patterns and inference. The machine learning models 71520 may build one or more mathematical models based on training data to make predictions and/or decisions without being explicitly programmed to perform the

specific tasks. In example implementations, machine learning models may perform classification, regression, clustering, anomaly detection, recommendation generation, digital twin creation and/or other tasks.

[2742] In example embodiments, the machine learning models 71520 may perform various types of classification based on the input data. Classification is a predictive modeling problem where a class label is predicted for a given example of input data. For example, the machine learning models 71520 can perform binary classification, multi-class or multi-label classification. In example embodiments, the machine-learning model may output "confidence scores" that are indicative of a respective confidence associated with classification of the input into the respective class. In example embodiments, the confidence scores can be compared to one or more thresholds to render a discrete categorical prediction. In example embodiments, a certain number of classes (e.g., one) with the relatively largest confidence scores can be selected to render a discrete categorical prediction.

[2743] In example embodiments, the machine learning models 71520 may output a probabilistic classification. For example, the machine learning models 71520 may predict, given a sample input, a probability distribution over a set of classes. Thus, rather than outputting only the most likely class to which the sample input should belong, the machine learning models 71520 can output, for each class, a probability that the sample input belongs to such class. In example embodiments, the probability distribution over all possible classes can sum to one. In example embodiments, a Softmax function, or other type of function or layer can be used to turn a set of real values respectively associated with the possible classes to a set of real values in the range (0, 1) that sum to one. In example embodiments, the probabilities provided by the probability distribution can be compared to one or more thresholds to render a discrete categorical prediction. In example embodiments, only a certain number of classes (e.g., one) with the relatively largest predicted probability can be selected to render a discrete categorical prediction.

[2744] In example embodiments, the machine learning models 71520 may perform regression modeling and related processes to provide output data in the form of a continuous numeric value. As examples, the machine learning models 71520 may perform linear regression, polynomial regression, logistic regression, nonlinear regression, or some other modeling process. As described, in example embodiments, a Softmax function or other function or layer can be used to squash a set of real values respectively associated with a two or more possible classes to a set of real values in the range (0, 1) that sum to one. For example, the machine learning models 71520 can perform linear regression, polynomial regression, or nonlinear regression. As examples, the machine learning models 71520 can perform simple regression or multiple regression. As described above, in some implementations, a Softmax function or other function or layer can be used to squash a set of real values respectively associated with a two or more possible classes to a set of real values in the range (0, 1) that sum to one.

[2745] In example embodiments, the machine learning models 71520 may perform various types of clustering. For example, the machine learning models 71520 may identify one or more previously-defined clusters to which the input data most likely corresponds. In some

implementations in which the machine learning models 71520 perform clustering, the machine learning models 71520 can be trained using unsupervised learning techniques.

[2746] In example embodiments, the machine learning models 71520 may perform anomaly detection or outlier detection. For example, the machine learning models 71520 can identify input data that does not conform to an expected pattern or other characteristic (e.g., as previously observed from previous input data). As examples, the anomaly detection can be used for fraud detection or system failure detection.

[2747] In some implementations, the machine learning models 71520 may provide output data in the form of one or more recommendations. For example, the machine learning models 71520 may be included in a recommendation system or engine. As an example, given input data that describes previous outcomes for certain entities (e.g., a score, ranking, or rating indicative of an amount of success or enjoyment), the machine learning models 71520 may output a suggestion or recommendation of one or more additional entities that, based on the previous outcomes, are expected to have a desired outcome.

[2748] As described above, the machine learning models 71520 may be or may include one or more of various different types of machine-learned models. Examples of such different types of machine-learned models are provided below for illustration. One or more of the example models described below can be used (e.g., combined) to provide the output data in response to the input data. Additional models beyond the example models provided herein can be used as well.

[2749] In some implementations, the machine learning models 71520 may be or may include one or more classifier models such as, for example, linear classification models; quadratic classification models; and the like. The machine learning models 71520 may be or may include one or more regression models such as, for example, simple linear regression models; multiple linear regression models; logistic regression models; stepwise regression models; multivariate adaptive regression splines; locally estimated scatterplot smoothing models; and the like.

[2750] In some examples, the machine learning models 71520 may be or may include one or more decision tree-based models such as, for example, classification and/or regression trees; chi-squared automatic interaction detection decision trees; decision stumps; conditional decision trees; and the like.

[2751] The machine learning models 71520 may be or may include one or more kernel machines. In some implementations, the machine learning models 71520 may be or may include one or more support vector machines. The machine learning models 71520 may be or may include one or more instance-based learning models such as, for example, learning vector quantization models; self-organizing map models; locally weighted learning models; and the like. In some implementations, the machine learning models 71520 may be or may include one or more nearest neighbor models such as, for example, k-nearest neighbor classifications models; k-nearest neighbors regression models; and the like. The machine learning models 71520 may be or may include one or more Bayesian models such as, for example, naïve Bayes models; Gaussian naïve Bayes models; multinomial naïve Bayes models; averaged one-dependence estimators; Bayesian networks; Bayesian belief networks; hidden Markov models; and the like.

[2752] In some implementations, the machine learning models 71520 may be or may include one or more artificial neural networks (also referred to simply as neural networks). A neural network may include a group of connected nodes, which also can be referred to as neurons or perceptrons. A neural network may be organized into one or more layers. Neural networks that include multiple layers may be referred to as "deep" networks. A deep network may include an input layer, an output layer, and one or more hidden layers positioned between the input layer and the output layer. The nodes of the neural network may be connected or non-fully connected.

[2753] The machine learning models 71520 may be or may include one or more feed forward neural networks. In feed forward networks, the connections between nodes do not form a cycle. For example, each connection can connect a node from an earlier layer to a node from a later layer.

[2754] In some instances, the machine learning models 71520 may be or may include one or more recurrent neural networks. In some instances, at least some of the nodes of a recurrent neural network can form a cycle. Recurrent neural networks can be especially useful for processing input data that is sequential in nature. In particular, in some instances, a recurrent neural network may pass or retain information from a previous portion of the input data sequence to a subsequent portion of the input data sequence through the use of recurrent or directed cyclical node connections.

[2755] In some examples, sequential input data may include time-series data (e.g., sensor data versus time or imagery captured at different times). For example, a recurrent neural network may analyze sensor data versus time to detect or predict a swipe direction, to perform handwriting recognition, etc. Sequential input data may include words in a sentence (e.g., for natural language processing, speech detection or processing, and the like); notes in a musical composition; sequential actions taken by a user (e.g., to detect or predict sequential application usage); sequential object states; and the like.

[2756] Example recurrent neural networks include long short-term (LSTM) recurrent neural networks; gated recurrent units; bi-direction recurrent neural networks; continuous time recurrent neural networks; neural history compressors; echo state networks; Elman networks; Jordan networks; recursive neural networks; Hopfield networks; fully recurrent networks; sequence-to-sequence configurations; and the like.

[2757] In some examples, the machine learning models 71520 may be or may include one or more non-recurrent sequence-to-sequence models based on self-attention, such as Transformer networks.

[2758] In some implementations, the machine learning models 71520 may be or may include one or more convolutional neural networks. In some instances, a convolutional neural network may include one or more convolutional layers that perform convolutions over input data using learned filters.

[2759] Filters may also be referred to as kernels. Convolutional neural networks may be especially useful for vision problems such as when the input data includes imagery such as still images or video. However, convolutional neural networks may also be applied for natural language processing.

[2760] In some examples, the machine learning models 71520 may be or may include one or more generative networks such as, for example, generative adversarial networks. Generative networks may be used to generate new data such as new images or other content.

[2761] The machine learning models 71520 may be or may include an autoencoder. In some instances, the aim of an autoencoder may learn a representation (e.g., a lower-dimensional encoding) for a set of data, typically for the purpose of dimensionality reduction. For example, in some instances, an autoencoder may seek to encode the input data and then provide output data that reconstructs the input data from the encoding. Recently, the autoencoder concept has become more widely used for learning generative models of data. In some instances, the autoencoder may include additional losses beyond reconstructing the input data.

[2762] The machine learning models 71520 may be or may include one or more other forms of artificial neural networks such as, for example, deep Boltzmann machines; deep belief networks; stacked autoencoders; and the like. Any of the neural networks described herein may be combined (e.g., stacked) to form more complex networks.

[2763] The machine learning models 71520 may include one or more clustering models such as, for example, k-means clustering models; k-medians clustering models; expectation maximization models; hierarchical clustering models; and the like.

[2764] In some implementations, the machine learning models 71520 may perform one or more dimensionality reduction techniques such as, for example, principal component analysis; kernel principal component analysis; graph-based kernel principal component analysis; principal component regression; partial least squares regression; Sammon mapping; multidimensional scaling; projection pursuit; linear discriminant analysis; mixture discriminant analysis; quadratic discriminant analysis; generalized discriminant analysis; flexible discriminant analysis; autoencoding; and the like.

[2765] In some implementations, the machine learning models 71520 may perform or be subjected to one or more reinforcement learning techniques such as Markov decision processes; dynamic programming; Q functions or Q-learning; value function approaches; deep Q-networks; differentiable neural computers; asynchronous advantage actor-critics; deterministic policy gradient; and the like.

Neural Networks for Machine Learning and Artificial Intelligence

[2766] In embodiments of the present disclosure, including embodiments involving artificial intelligence, machine learning, automation (including robotic process automation, remote control, autonomous operation, automated configuration, and the like), expert systems, self-organization, adaptive intelligent systems for prediction, classification, optimization, and the like, may benefit from the use of a neural network, such as a neural network trained for pattern recognition, for classification of one or more parameters, characteristics, or phenomena, for support of autonomous control, and other purposes.

[2767] Neural networks (or artificial neural networks) are a family of statistical learning models inspired by biological neural networks and are used to estimate or approximate functions that may depend on a large number of inputs and are generally unknown. Neural networks represent a system

of interconnected "neurons" which send messages to each other. The connections have numeric weights that can be tuned based on experience, making neural nets adaptive to inputs and capable of learning.

[2768] References to artificial intelligence, neural networks or neural net throughout this disclosure should be understood to encompass a wide range of different types of machine learning systems, neural networks, such as feed forward neural networks, convolutional neural networks (CNN), recurrent neural networks (RNN), long short-term memory (LSTM) neural networks, gated recurrent unit (GRU) neural networks, self-organizing map (SOM) neural networks (e.g., Kohonen self-organizing neural networks), autoencoder (AE) neural networks, encoder-decoder neural networks, modular neural networks, or variations, hybrids or combinations of the foregoing, or combinations with reinforcement learning (RL) systems or other expert systems, such as rule-based systems, and model-based systems (including ones based on physical models, statistical models, flow-based models, biological models, biomimetic models and the like).

[2769] The neural networks referenced and described herein may have a variety of nodes or neurons, which may perform a variety of functions on inputs, such as inputs received from sensors or other data sources, including other nodes to predict one or more outputs. Functions may involve weights, features, feature vectors, and the like. Neurons may include perceptrons, neurons that mimic biological functions (such as the human senses of touch, vision, taste, hearing, and smell), and the like. Neural networks can employ multiple layers of operations including one or more hidden layers situated between an input layer and an output layer. The output of each layer can be used as input to another layer, e.g., the next hidden layer or the output layer. The output of a particular neuron can be a weighted sum of the inputs to the neuron, adjusted with a bias and multiplied by an activation function, e.g., a rectified linear unit (ReLU) or a sigmoid function.

[2770] In many embodiments, an expert system or neural network may be trained, such as by a human operator or supervisor, or based on a data set, model, or the like. Training a neural network can involve providing inputs to the untrained neural network to generate predicted outputs, comparing the predicted outputs to expected outputs, and updating the algorithm's weights and biases to account for the difference between the predicted outputs and the expected outputs. Specifically, a cost function can be used to calculate a difference between the predicted outputs and the expected outputs. By computing the derivative of the cost function with respect to the weights and biases of the network, the weights and biases can be iteratively adjusted over multiple cycles to minimize the cost function. Training may be complete when the predicted outputs satisfy a convergence condition (e.g., a small magnitude of calculated cost as determined by the cost function).

[2771] Training may include presenting the neural network with one or more training data sets that represent values (including the many types described throughout this disclosure), as well as one or more indicators of an outcome, such as an outcome of a process, an outcome of a calculation, an outcome of an event, an outcome of an activity, or the like. Training may include training in optimization, such as training a neural network to optimize one or more systems based on one or more optimization approaches, such as Bayesian approaches, parametric Bayes classifier

approaches, k-nearest-neighbor classifier approaches, iterative approaches, interpolation approaches, Pareto optimization approaches, algorithmic approaches, and the like. Feedback may be provided in a process of variation and selection, such as with a genetic algorithm that evolves one or more solutions based on feedback through a series of rounds.

[2772] In example embodiments, a plurality of neural networks may be deployed in a cloud platform that receives data streams and other inputs collected (such as by mobile data collectors) in one or more environments and transmitted to the cloud platform over one or more networks, including using network coding to provide efficient transmission. In the cloud platform, optionally using massively parallel computational capability, a plurality of different neural networks of various types (including modular forms, structure-adaptive forms, hybrids, and the like) may be used to undertake prediction, classification, control functions, and provide other outputs as described in connection with expert systems disclosed throughout this disclosure. The different neural networks may be structured to compete with each other (optionally including use evolutionary algorithms, genetic algorithms, or the like), such that an appropriate type of neural network, with appropriate input sets, weights, node types and functions, and the like, may be selected, such as by an expert system, for a specific task involved in a given context, workflow, environment process system, or the like.

[2773] In example embodiments, methods and systems described herein that involve an expert system or self-organization capability may use a feed forward neural network, which moves information in one direction, such as from a data input, like a source of data about an individual, through a series of neurons or nodes, to an output. Data may move from the input nodes to the output nodes, optionally passing through one or more hidden nodes, without loops. In example embodiments, feed forward neural networks may be constructed with various types of units, such as binary McCulloch-Pitts neurons, the simplest of which is a perceptron.

[2774] FIG. 264 depicts a diagrammatic view illustrating an example implementation of a data processing system using a neural network to provide real-time, adaptive control of a dynamic vision system including object classification and determination of object position, orientation and motion, according to some embodiments of the present disclosure.

[2775] Neural networks generally comprise an interconnected group of nodes organized into multiple layers of nodes. For example, the neural network architecture may comprise at least an input layer, one or more hidden layers, and an output layer with each layer comprising a plurality of nodes or neurons that respond to different combinations of inputs from the previous layers. The input for the input layer is received directly from sensors and imaging data whereas the hidden layers use output of nodes in previous layers as their input. The connections between the neurons have numeric weights that determine how much relative effect an input has on the output value of the node in question. The neural network may comprise any total number of layers, and any number of hidden layers, where the hidden layers function as trainable feature extractors that allow mapping of a set of input data to a preferred output value or set of output values. The output layer provides an output in the form of one or more predictions, recommendations, classifications,

optimizations or decisions related to one or more components or parameters of the dynamic vision system 71300.

[2776] Referring to FIG. 6, the input layer may comprise one or more real-time data streams of imaging data that provide an indication of the current state of the object. This data may be fed to the neural network, which, in example embodiments, has been previously trained using one or more training data sets. The input layer may include a plurality of input nodes 71602, 71604, 71606, 71608 and 71610 that may provide input data (e.g., sensor data, image data, audio data, etc.) to the neural network 71100. The input data may be from different sources and may include library data x1, simulation data x2, user input data x3, training data x4 and outcome data x6. The input nodes 71602, 71604, 71606 71608 and 71610 may pass on the information to the next layer, and, in example embodiments, no computation may be performed by the input nodes. Hidden layer(s) may include a plurality of nodes, such as nodes 71612, 71614, and 71616 that may process the information from the input layer based on the weights of the connections between the input layer and the hidden layer and transfer information to the output layer. The output layer may include an output node 71618 which processes information based on the weights of the connections between the hidden layer and the output layer and is responsible for computing and transferring information from the network to the outside world, such as optimizing an optical parameter, classifying certain objects or defects, or predicting a condition or an action.

[2777] In example embodiments, the neural network 71600 may include two or more hidden layers and may be referred to as a deep neural network. In example embodiments, the layers may be constructed so that the first layer detects a set of primitive patterns in the input (e.g., image) data, the second layer detects patterns of patterns and the third layer detects patterns of those patterns. In example embodiments, a node in the neural network 71600 may have connections to all nodes in the immediately preceding layer and the immediate next layer. Thus, the layers may be referred to as fully connected layers. In example embodiments, a node in the neural network 71600 may have connections to only some of the nodes in the immediately preceding layer and the immediate next layer. Thus, the layers may be referred to as sparsely connected layers.

[2778] Each neuron in the neural network consists of a weighted combination (e.g., linear combination) of its inputs, and the computation on each neural network layer may be described as a multiplication of an input matrix and a weight matrix. A bias matrix may then be added to the resulting product matrix to account for the threshold of each neuron in the next level. Further, an activation function may be applied to each resultant value, and the resulting values may be placed in the matrix for the next layer. Thus, the output from a node in the neural network may be represented as:

[2779] $yi = f(\sum xiwi + bi)$

[2780] where f is the activation function, $\sum xiwi$ is the weighted sum of input matrix and bi is the bias matrix.

[2781] The activation function determines the activity level or excitation level generated in the node as a result of an input signal of a particular size. The purpose of the activation function is to introduce non-linearity into the output of a neural network node because most real-world functions

are non-linear and it is desirable that the neurons can learn these non-linear representations. Several activation functions may be used in an artificial neural network. One example activation function is the sigmoid function σ(x), which is a continuous S-shaped monotonically increasing function that asymptotically approaches fixed values as the input approaches plus or minus infinity. The sigmoid function σ(x) takes a real-valued input and transforms it into a value between 0 and 1:

[2782] $\sigma(x)=1/(1+\exp(-x))$.

[2783] Another example activation function is the tanh function, which takes a real-valued input and transforms it into a value within the range of [−1, 1]:

[2784] $\tanh(x)=2\sigma(2x)-1$

[2785] A third example activation function is the rectified linear unit (ReLU) function. The ReLU function takes a real-valued input and thresholds it above zero (i.e., replacing negative values with zero):

[2786] $f(x)=\max(0, x)$.

[2787] In the example shown in FIG. 6, nodes 71602, 71604, 71606, 71608 and 71610 in the input layer may take external inputs x1, x2, x3, x4 and x6 which may be numerical values depending upon the input dataset. It will be understood that even though only five inputs are shown in FIG. 6, in various implementations, a node may include tens, hundreds, thousands, or more inputs. As discussed herein, no computation is performed on input layer and thus the outputs from nodes 71602, 71604, 71606, 71608 and 71610 of input layer are x1, x2, x3, x4 and x6 respectively, which are fed into the hidden layer. The output of node 71612 in the hidden layer may depend on the outputs from input layer (x1, x2, x3, x4 and x6) and weights associated with connections (w1, w2, w3, w4 and w6). Thus, the output from node 71612 may be computed as:

[2788] $Y71612=f(x1w1+x2w2+x3w3+x4w4+x6w6+b71612)$

[2789] The outputs from the nodes 71614 and 71616 in the hidden layer may also be computed in a similar manner and then be fed to the node 71618 in the output layer. Node 71618 in the output layer may perform similar computations (using weights v1, v2 and v3 associated with the connections) as the nodes 71612, 71614 and 71616 in the hidden layers.

[2790] $Y71618=f(y71612v1+y71614v2+y71616v3+b71618)$

[2791] Where Y71618 is the output of the neural network 71600.

Training

[2792] As described herein, the connections between nodes in a neural network have associated weights. Weights determine how much relative effect an input value has on the output value of the node in question. Before the network is trained, random values are selected for each of the weights. The weights are adjusted during the training process and this adjustment of weights to determine the best set of weights that maximize the accuracy of the neural network is referred to as training. For every input in a training dataset, the output of the artificial neural network may be observed and compared with the expected output, and the error between the expected output and the observed output may be propagated back to the previous layer. The weights may be adjusted accordingly based on the error. This process is repeated until the output error is below a predetermined threshold.

[2793] In example embodiments, backpropagation (e.g., backward propagation of errors) may be utilized with an optimization method such as gradient descent to adjust weights and update the neural network characteristics. Backpropagation may be a supervised training scheme that learns from labeled training data and errors at the nodes by changing parameters of the neural network to reduce the errors. For example, a result of forward propagation (e.g., output activation value(s)) determined using training input data is compared against a corresponding known reference output data to calculate a loss function gradient. The gradient may be then utilized in an optimization method to determine new updated weights in an attempt to minimize a loss function. For example, to measure error, the mean square error is determined using the equation:

[2794] E=(target−output)2

[2795] To determine the gradient for a weight "w," a partial derivative of the error with respect to the weight may be determined, where:

[2796] gradient=$\partial E/\partial w$

[2797] The calculation of the partial derivative of the errors with respect to the weights may flow backwards through the node levels of the neural network. Then a portion (e.g., ratio, percentage, etc.) of the gradient is subtracted from the weight to determine the updated weight. The portion may be specified as a learning rate "a." Thus an example equation of determining the updated weight is given by the formula:

[2798] w new =w old $-\alpha\partial E/\partial w$

[2799] The learning rate may be selected such that it is not too small (e.g., a rate that is too small may lead to a slow convergence to the desired weights) and not too large (e.g., a rate that is too large may cause the weights to not converge to the desired weights). After the weight adjustment, the network should perform better than before for the same input because the weights have now been adjusted to minimize the errors.

[2800] In example embodiments, a neural network model may be used directly to determine adjustments to optical parameters using training or learning of a neural network model. Initially, the model may be allowed to choose randomly from a range of values for each input optical control parameter or action. If the sequence of optical control parameter adjustments or actions leads to an incorrect prediction/classification, it may be scored as leading to an undesirable (or negative) outcome. Repetition of the process using different sets of randomly chosen values for each optical control parameter or action leads to reinforcement of those sequences that least to desirable (or positive) outcomes. Ultimately, the neural network model "learns" what adjustments to make to a set or sequence of optical control parameters or actions in order to achieve the target outcome, i.e., a correct prediction or classification.

[2801] In example embodiments, methods and systems described herein may use a convolutional neural network (referred to in some cases as a CNN, a ConvNet, a shift invariant neural network, or a space invariant neural network), wherein the units are connected in a pattern similar to the visual cortex of the human brain.

[2802] FIG. 265 depicts a diagrammatic view illustrating an example implementation of the processing system using a convolutional neural network (CNN) to provide automatic classification of objects by the dynamic vision system according to some embodiments of the present disclosure.

[2803] A convolutional neural network (CNN) is a specialized neural network for processing data having a known, grid-like topology, such as image data. Accordingly, CNNs are commonly used for classification, object recognition and machine vision applications, but they also may be used for other types of pattern recognition such as speech and language processing.

[2804] A convolutional neural network learns highly non-linear mappings by interconnecting layers of artificial neurons arranged in many different layers with activation functions that make the layers dependent. It may include one or more convolutional layers, interspersed with one or more sub-sampling layers and non-linear layers, which are typically followed by one or more fully connected layers.

[2805] Referring to FIGS. 3 and 7, a CNN may include an input layer with an input object concept 71720 to be classified by the CNN, a hidden layer which in turn may include one or more convolutional layers, interspersed with one or more activation or non-linear layers (e.g., ReLU) and pooling or sub-sampling layers and an output layer- typically including one or more fully connected layers. The input object concept 71720 may be represented by a matrix of pixels and may denote an input data stream generated at the sensors 71310 by the liquid lens 71308.

[2806] As shown in FIG 7, the input image 71720 may be processed by the hidden layer, which includes sets of convolutional and activation layers 71722 and 71726, each followed by pooling layers 71724 and 71728.

[2807] The convolutional layers of the convolutional neural network may serve as feature extractors capable of learning and decomposing the input object concept into hierarchical features. The convolution layers may perform convolution operations on the input concept where a filter (also referred as a kernel or feature detector) may slide over the input object concept at a certain step size (referred to as the stride). For every position (or step), element-wise multiplications between the filter matrix and the overlapped matrix in the input object concept may be calculated and summed to get a final value that represents a single element of an output matrix constituting a feature map. The feature map refers to data that represents various features of the input object concept data. The activation or non-linear layers use different non-linear trigger functions to signal distinct identification of likely features on each hidden layer. Non-linear layers may use a variety of specific functions to implement the non-linear triggering, including the rectified linear units (ReLUs), hyperbolic tangent, absolute of hyperbolic tangent and sigmoid functions. In one implementation, a ReLU activation implements the function y=max(x, 0) and keeps the input and output sizes of a layer the same. One advantage of using ReLU is that the convolutional neural network is trained many times faster. ReLU is a non-continuous, non-saturating activation function that is linear with respect to the input if the input values are larger than zero and zero otherwise.

[2808] As shown in FIG. 7, the first convolution and activation layer 71722 may perform convolutions on the input object concept 71720 using multiple filters followed by non-linearity operation (e.g., ReLU) to generate multiple output matrices (or feature maps) 71730. The number

of filters used may be referred to as the depth of the convolution layer. Thus, the first convolution and activation layer 71722 in the example of FIG. 265 has a depth of three and generates three feature maps using three filters. Feature maps 71730 may then be passed to the first pooling layer 71724 that may sub-sample or down-sample the feature maps using a pooling function to generate an output matrix 71732. The pooling function replaces the feature map with a summary statistic to reduce the spatial dimensions of the extracted feature map thereby reducing the number of parameters and computations in the network. Thus, the pooling layer reduces the dimensionality of the feature maps while retaining the most important information. The pooling function can also be used to introduce translation invariance into the neural network, such that small translations to the input do not change the pooled outputs. Different pooling functions may be used in the pooling layer, including max pooling, average pooling, and 12-norm pooling.

[2809] The output matrix 71732 may then be processed by a second convolution and an activation layer 71726 to perform convolutions and non-linear activation operations (e.g., ReLU) as described above to generate feature maps 71734. In the example shown in FIG. 7, a second convolution and the activation layer 71726 may have a depth of five. The feature maps 71734 may then be passed to a pooling layer 71728, where the feature maps 71734 may be subsampled or down-sampled to generate an output matrix 71736.

[2810] The output matrix 71736 generated by the pooling layer 71728 may then be processed by one or more fully connected layer 71738 that forms a part of the output layer. The fully connected layer 71738 has a full connection with all the feature maps of the output matrix 71736 of the pooling layer 71728. In example embodiments, the fully connected layer 71738 may take the output matrix 71736 generated by the pooling layer 71728 as the input in vector form and perform high-level determination to output a feature vector containing information of the structures in the input object concept 71720. In embodiments the fully connected layer 71738 may classify the object in input object concept 71720 into one of several categories, such as using a Softmax function. In example embodiments, the Softmax function may be used as the activation function in the output layer and may take a vector of real-valued scores and map it to a vector of values between zero and one that sum to one. In example embodiments, other classifiers, such as a support vector machine (SVM) classifier, may be used.

[2811] In example embodiments, one or more normalization layers may be added to the CNN to normalize the output of the convolution filters. The normalization layer may provide whitening or lateral inhibition, avoid vanishing or exploding gradients, stabilize training, and enable learning with higher rates and faster convergence. In example embodiments, the normalization layers may be added after the convolution layer but before the activation layer.

[2812] A CNN may thus be seen as multiple sets of convolution, activation, pooling, normalization and fully connected layers stacked together to learn, enhance and extract implicit features and patterns in the input object concept. A layer, as used herein, can refer to one or more components that operate with similar function by mathematical or other functional means to process received inputs to generate/derive outputs for a next layer with one or more other components for further processing within the CNN.

[2813] The initial layers of the CNN (e.g., convolution layers), may extract low level features such as edges and/or gradients from the input object concept 71720. Subsequent layers may extract or detect progressively more complex features and patterns such as presence of curvatures and textures in image data and so on. The output of each layer may serve as an input of a succeeding layer in the CNN to learn hierarchical feature representations from data in the input object concept 71720. This allows convolutional neural networks to efficiently learn increasingly complex and abstract visual concepts.

[2814] Although only two convolution layers are shown in the example, the present disclosure is not limited to the example architecture, and the CNN architecture may comprise any number of layers in total, and any number of layers for convolution, activation and pooling. In example embodiments, a convolutional neural network may be deployed with a large number of neurons (e.g., 100,000, 600,000 or more), with multiple (e.g., 10, 60, 100 or more) layers, and with thousands of parameters. For example, there have been many variations and improvements over the basic CNN model described above. Some examples include Alexnet, GoogLeNet, VGGNet (that stacks many layers containing narrow convolutional layers followed by max pooling layers), Residual network or ResNet (that uses residual blocks and skip connections to learn residual mapping), DenseNet (that connects each layer of CNN to every other layer in a feed-forward fashion), Squeeze and excitation networks (that incorporate global context into features) and AmobeaNet (that uses evolutionary algorithms to search and find optimal architecture for image recognition).

Training of Convolutional Neural Networks

[2815] The training process of a convolutional neural network may be similar to the training process discussed in FIG. 264 with respect to the neural network 71600. First, parameters and weights (including the weights in the filters and weights for the fully connected layer) may be assigned, such as randomly assigned. Then, during training, a set of training images/object concepts (in the case of CNNs used for object recognition) in which the objects have been detected and classified is provided as the input to the CNN, which performs the forward propagation steps. In other words, the CNN applies convolution, non-linear activation, and pooling layers to each training image/object concept to determine the classification vectors (i.e., to detect and classify each training image/object concept). These classification vectors may be compared with the predetermined classification vectors. The error (e.g., the squared sum of differences, log loss, Softmax log loss) between the classification vectors of the CNN and the predetermined classification vectors may be determined. This error is then employed to update the weights and parameters of the CNN in a backpropagation process which may use gradient descent and may include one or more iterations. The training process may be repeated for each training image/object concept in the training set.

[2816] The training process and inference process described herein may be performed on hardware, software, or a combination of hardware and software. However, training a convolutional neural network or using the trained CNN for inference generally requires a significant amount of computation power to perform, for example, the matrix multiplications or convolutions. Thus,

specialized hardware circuits, such as graphic processing units (GPUs), tensor processing units (TPUs), neural network processing units (NPUs), FPGAs, ASICs, or other highly parallel processing circuits may be used for training and/or inference. Training and inference may be performed on a cloud, on a data center, or on a device.

[2817] In example embodiments, capsule networks may be employed to use fewer labeled training examples to achieve similar classification performance of CNNs.

[2818] In example embodiments, transformer-based, encoder-decoder architectures using attention mechanisms may be used in conjunction with or in place of convolutional neural networks.

[2819] FIG. 266 depicts an example embodiment of a transformer neural network 71800 used in conjunction with the dynamic vision system 71300. The transformer neural network 71800, in an embodiment, may include three input stages and five output stages to transform an input sequence into an output sequence. The example transformer includes an encoder 71802 and a decoder 71804. The encoder 71802 processes input, and the decoder 71804 generates output probabilities, for example. The encoder 71802 may include three stages, and the decoder 71804 may include five stages. Encoder 71802 stage 1 represents an input as a sequence of positional encodings added to embedded inputs. Encoder 71802 stages 2 and 3 include N layers (e.g., N=6, etc.) in which each layer includes a position-wise feedforward neural network (FNN) and an attention-based sublayer. Each attention-based sublayer of encoder 71802 stage 2 includes four linear projections and multi-head attention logic to be added and normalized to be provided to the position-wise FNN of encoder 71802 stage 3. Encoder 71802 stages 2 and 3 employ a residual connection followed by a normalization layer at their output.

[2820] The example decoder 71804 processes an output embedding as its input with the output embedding shifted right by one position to help ensure that a prediction for position i is dependent on positions previous to/less than i. In stage 2 of the decoder 71804, masked multi-head attention is modified to prevent positions to attend to subsequent positions. Stages 3-4 of the decoder 71804 include N layers (e.g., N=6, etc.) in which each layer includes a position-wise FNN and two attention-based sublayers. Each attention-based sublayer of decoder 71404 stage 3 includes four linear projections and multi-head attention logic to be added and normalized to be provided to the position-wise FNN of decoder 71804 stage 4. Decoder 71804 stages 2-4 employ a residual connection followed by a normalization layer at their output. Decoder 71404 stage 6 provides a linear transformation followed by a softmax function to normalize a resulting vector of K numbers into a probability distribution 71806 including K probabilities proportional to exponentials of the K input numbers.

[2821] FIG. 267 depicts a schematic view illustrating an example implementation of a dynamic vision system depicting a detailed view of various components along with integration of the dynamic vision system with one or more third party systems according to some embodiments of the present disclosure. The dynamic vision system 71900 may include a liquid lens optical assembly 71304 configured to capture data from various data sources 71902 including vision sensors 71904, feedback sources 71906 providing outcome data from the machine learning system,

environment control 71908 generating data in response to a change in environment factors including temperature, pressure, humidity, vibrations etc., lighting control 71910 generating data in response to a change in source lighting including color, color temperature, timing (PWM), amplitude etc. and data library 71912.

[2822] The data storage and management system 71914 may maintain a record of state and event data for various components and subsystems of the dynamic vision system 71300 such that any of the services, applications, programs, or the like may access a common data source (which may comprise a single logical data source that is distributed across disparate physical and/or virtual storage locations). The data storage and management system 71914 may include a memory subsystem for storage of instructions and data and a file storage subsystem providing persistent storage for program and data files. Further, the storage and management system 71914 may include capabilities such as data allocation, data caching, data pruning and data management and access to and control of intelligence and data resources.

[2823] The processing system 71306 may process the data captured by liquid lens optical assembly 71304 and stored in data storage and management system 71914 to optimize and adjust the optical parameters in real time through the machine learning system 71324 and the artificial intelligence system 71326, the digital twin system 71320 and the control system 71314 as described in detail in FIGS. 2, 3 4 and 5, or elsewhere herein.

[2824] In example embodiments, a set of applications 71916 may enable the dynamic vision system 71300 to present meaning information to a user and enable the user to perform specific vision tasks. Some examples of applications provided on the dynamic vision system 71300 include particle filter 71918, 3D model generation 71920, Location or motion prediction 71922, Visual SLAM 71924, defect detection 71926 and adversarial neural network detection 71928.

[2825] In example embodiments, the dynamic vision system 71300 may integrate with one or more third party systems 71930 through connectivity facilities including interfaces, network connections, ports, application programming interfaces (APIs), brokers, services, connectors, wrappers, containers, wired or wireless communication links, human-accessible interfaces, software interfaces, micro-services, SaaS interfaces, PaaS interfaces, IaaS interfaces, cloud capabilities, or the like. The connectivity facilities may facilitate the transfer of data between the dynamic vision system 71300 and the one or more third party systems 71930.

[2826] In example embodiments, the dynamic vision system 71300 may integrate into or with a set of value chain network entities for quality control inspections and sorting objects in a production assembly line or logistics chain wherein the liquid lens is configured to quickly adjust focus to accommodate for, recognize and sort objects located at various working distances or objects of different heights.

[2827] In example embodiments, the dynamic vision system 71300 may integrate into or with a set of autonomous vehicle systems to scan the vehicle environment and monitor the distance between the vehicle from other objects on the road.

[2828] In example embodiments, the dynamic vision system 71300 may integrate into or with an interactive head-mounted device configured to display virtual content with an electrically

adjustable liquid lens for providing a correction for the displayed content by adjusting the electrically adjustable liquid lens.

[2829] In example embodiments, the dynamic vision system 71300 may integrate into or with an unmanned automotive vehicle (UAV) navigation system to help control the position or course of the UAV in three dimensions.

[2830] Some non-limiting examples of third party systems 71930 that may integrate with dynamic vision system 71300 for incorporating vision capability include IoT system 71932, value chain network system 71934, manufacturing execution system (MES) 71936, robot/cobot system 71938, automotive system 71940, 3D printing system, ophthalmic system, surgical system, microscopy system, exoskeleton system, prosthetics system, biometrics system, quality management system (QMS), compliance system, certification system, and the like.

[2831] In example embodiments, the integration of the dynamic vision system 71300 with the more third-party systems 71930 takes into account the specific needs and requirements of the third party systems 71930 and may customize certain components and applications of the dynamic vision system 71300 based on such requirements. For example, when integrating with a 3D printing system, defect detection may be provided whereas integration with a robotic cleaning system may benefit from the inclusion of virtual SLAM 71924.

Use Cases

[2832] In example embodiments, the dynamic vision system 71300 may be configured to be a component of counting devices which may be placed strategically throughout a retail store. The system 71300 may gather data about where customers spend their time and for how long. The insights derived by such customer analytics may improve retail stores' understanding of consumer interaction and help improve store layout optimization and staffing levels.

[2833] In example embodiments, the dynamic vision system 71300 may help retail stores improve customer experience and ensure social distancing through a footfall counting system to measure the number of people that pass through a certain passage or entrance. Footfall counting may also form the basis of retail analytics, queue management and space utilization applications for the retail store.

[2834] In example embodiments, the dynamic vision system 71300 may be a component in the productivity analytics suite of applications to help track employee productivity thereby providing insights into employee time management and workplace collaboration in the retail environment.

[2835] Transportation

[2836] In example embodiments, the dynamic vision system 71300 may be configured to be a component of in-vehicle driver monitoring technologies helping monitor the physiological, psychological, emotional, physical and positional states of a driver and determine attentiveness through recognizing facial expressions, posture, gaze, and the like. Similar technologies may also be employed for passenger monitoring in autonomous or semi-autonomous vehicles.

[2837] In example embodiments, the dynamic vision system 71300 may be configured to be a component of an advanced driver assistance system (ADAS) for lane detection, pedestrian detection, traffic sign detection, collision avoidance and parking occupancy detection.

[2838] In example embodiments, the dynamic vision system 71300 may help monitor the physical and functional conditions of the roads and other infrastructure and well as determine traffic conditions thereby helping to identify the most appropriate routes for vehicle navigation.

[2839] In example embodiments, the dynamic vision system 71300 may form part of automotive cameras to scan the vehicle environment and monitor the distance between the vehicle from other objects on the road.

[2840] In example embodiments, the dynamic vision system 71300 may help an unmanned aerial vehicle (UAV) navigation system control the position or course of the UAV in three dimensions.

[2841] In example embodiments, the dynamic vision system 71300 may be used in manufacturing environments for quality control inspections and sorting of objects in a production assembly line or logistics chain. The conformable liquid lens based optical assembly may be configured to quickly adjust focus to accommodate objects located at various working distances or objects of different heights.

[2842] In example embodiments, the dynamic vision system 71300 may be used in manufacturing environments for productivity analytics to help track employee productivity thereby providing insights into employee time management and workplace collaboration.

[2843] In example embodiments, the dynamic vision system 71300 may be configured to monitor crops including identifying the crops that are ready to be harvested and detecting pests, weeds, and any nutritional deficiencies, assessing crop yields and testing for agricultural product quality.

[2844] In example embodiments, the dynamic vision system 71300 may be configured to monitor animals or livestock in farming, where livestock can be monitored remotely for security from predators, disease detection, changes in behavior and the like.

[2845] Apart from the various robotic vision use cases described in examples, the dynamic vision system 71300 may help robotic vision with motion planning and in identifying an optimal collision free path in real-time in a 3D workspace while taking into account various kinematic, geometric, physical and temporal constraints. Additionally, the dynamic vision system may help with identifying any moving obstacles and predicting the trajectory of the moving obstacle in the environment, which information may be considered for motion planning.

Anomaly Detection and Diagnosis

[2846] In various embodiments, a method of detecting an anomaly associated with a machine includes recording a data set associated with the machine; determining, by a first machine learning model, a label associated with the data set; determining whether the label is to be reviewed; and responsive to determining that the label is to be reviewed, subjecting the data set and the label to a review, and updating the label based on the review.

[2847] In various embodiments, a system for detecting an anomaly associated with a machine, the system includes a processor and a memory storing instructions that, when executed by the processor, cause the system to record a data set associated with the machine, determine, by a first machine learning model, a label associated with the data set, determine whether the label is to be reviewed, and responsive to determining that the label is to be reviewed, subject the data set and the label to a review, and update the label based on the review.

[2848] In various embodiments, a non-transitory computer-readable storage medium storing instructions that, when executed by a processor, cause the processor to detect an anomaly associated with a machine by recording a data set associated with the machine, determining, by a first machine learning model, a label associated with the data set, determining whether the label is to be reviewed, and responsive to determining that the label is to be reviewed, subjecting the data set and the label to a review, and updating the label based on the review.

[2849] In various embodiments, a method of deploying an agent within an environment includes associating, with the agent, at least one action, at least one instinct that motivates the agent to act within the environment, and at least one goal associated with a task within the environment, and associating, with the environment, at least one environmental rule that triggers the agent to perform the at least one action.

[2850] In various embodiments, a system for deploying an agent within an environment, the system includes a processor and a memory storing instructions that, when executed by the processor, cause the system to associate, with the agent, at least one action, at least one instinct that motivates the agent to act within the environment, and at least one goal associated with a task within the environment, and associate, with the environment, at least one environmental rule that triggers the agent to perform the at least one action.

[2851] In various embodiments, a non-transitory computer-readable storage medium storing instructions that, when executed by a processor, cause the processor to deploy an agent within an environment by associating, with an agent, at least one action, at least one instinct that motivates the agent to act within an environment, and at least one goal associated with a task within the environment, and associating, with the environment, at least one environmental rule that triggers the agent to perform the at least one action.

[2852] In various embodiments, a method of presenting an analysis of a machine included in an industrial facility includes generating a digital twin of the machine; determining at least one property of the digital twin based on a simulation of an operation of the machine; and generating a presentation of the industrial facility that includes a visualization of the digital twin and a visual indicator of the at least one property of the digital twin.

[2853] In various embodiments, a system for presenting an analysis of a machine included in an industrial facility includes a processor and a memory storing instructions that, when executed by the processor, cause the system to generate a digital twin of the machine, determine at least one property of the digital twin based on a simulation of an operation of the machine, and generate a presentation of the industrial facility that includes a visualization of the digital twin and a visual indicator of the at least one property of the digital twin.

[2854] In various embodiments, a non-transitory computer-readable storage medium storing instructions that, when executed by a processor, cause the processor to present an analysis of a machine included in an industrial facility, by generating a digital twin of the machine; determining at least one property of the digital twin based on a simulation of an operation of the machine; and generating a presentation of the industrial facility that includes a visualization of the digital twin and a visual indicator of the at least one property of the digital twin.

[2855] Locations such as industrial facilities often include one or more machines. The machines may be mechanical in nature (e.g., levers, pulleys, flywheels, or gasoline-powered vehicles), electrical in nature (e.g., motors, rotors, stators, fans, or electric vehicles), electronic in nature (e.g., servers, workstations, or communication equipment), and/or of any other nature. The term "machine" as used herein includes, without limitation, machines with one or more moving parts and/or machines with one or more nonmoving parts. Within any such machine, anomalies can arise that jeopardize continued functioning of the machine. For example, in a mechanical machine, anomalous misalignment of a moving part can lead to undesired contact with another part, causing friction, jamming, wear, breakage, heat, noise, or physical strain. In an electrical machine, anomalous electrical malfunctions such as arcing can lead to short circuits, burnt components or wires, or a fire. In an electronic machine, anomalous logical flaws can lead to errors, causing data loss, failed processes, damaged components, or loss of power.

[2856] Anomalies can be detectable by monitoring signals from the machine. For example, in a mechanical machine, physical sensors can detect a vibration of a moving part, which can indicate various anomalies associated with the moving part. In an electrical machine, voltage or current sensors can detect deviations in the expected electrical properties of an electrical component, which can indicate various anomalies associated with the electrical component. In an electronic machine, logic monitoring processes can monitor data to detect unexpected values or changes, which can indicate various anomalies associated with the data. In these and other scenarios, the signals detected by a sensor may include an abundance of data by which the anomaly can be detected and diagnosed. Such detection and diagnosis are sometimes achieved through analysis performed by a human, including a subject-matter expert (SME). However, the volume of data and/or the number of monitored machines may be difficult and/or inefficient for the human to analyze comprehensively and rapidly, resulting in erroneous and/or delayed detection and diagnosis. Such errors and delays may result in damage to the machines that could have been avoided or mitigated by the application of more comprehensive and/or timely techniques.

[2857] Presented herein are techniques for anomaly detection and diagnosis. The techniques include machine learning models that can automatically detect anomalies associated with a machine (e.g., vibration data detected by vibration sensors). The techniques include recording a data set associated with the machine, determining, by a first machine learning model, a label associated with the data set, determining whether the label is to be reviewed, and responsive to determining that the label is to be reviewed, subjecting the data set and the label to a review, and updating the label based on the review.

[2858] Additionally, in the field of anomaly detection and diagnosis as well as other fields of science, technology, industry, and the like, machine learning is often applied by training one or more models (e.g., one or more artificial neural networks) on a data set. These models can be trained to perform one or more tasks. These models may also include, incorporate, and/or model various forms of hard-coded logic (e.g., representing equations used in the relevant technical domain such as physics, chemistry, or electronics and/or knowledge or intuition of subject-matter experts in the technical domain). However, such models can be difficult to maintain, adapt to

changes to the tasks for which the models are trained, or extend to address new tasks, particularly those with poorly defined subject matter.

[2859] Presented herein are techniques for deploying an agent within an environment. The techniques involve associating, with the agent, at least one action, at least one instinct that motivates the agent to act within the environment, and at least one goal associated with a task within the environment, and associating, with the environment, at least one environmental rule that triggers the agent to perform the at least one action.

[2860] Additionally, in the field of anomaly detection and diagnosis as well as other fields of science, technology, industry, and the like, machines and pieces of equipment are often monitored and evaluated based on data collected by sensors. For example, an industrial facility may include a set of machines, each performing one or more tasks in one or more contexts and according to certain configurations and operating characteristics. Sensors can monitor various features of each machine during operation, and problems can be diagnosed by comparing the values of the monitored features with thresholds that are associated with operating tolerances. For example, sensors can detect vibrations in machines that indicate the development or progression of a mechanical issue. However, monitoring many such values for the set of machines of the industrial facility may be difficult due to the volume of data and variance between machines. Further, while a particular value of a particular feature that is outside of a threshold might indicate an anomaly, determining the particular nature or cause of the anomaly can be difficult without an extended inspection of the machine. For example, when a sensor detects an occurrence of a vibration associated with a machine, it may be difficult to determine whether the vibration is caused by a misalignment of parts, wear or damage to bearings, insufficient lubrication between parts, or myriad other causes. Such determination might involve deactivating, disassembling, inspecting, and/or testing the machine, which can be time-consuming and difficult.

[2861] Presented herein are techniques for using digital twins in anomaly detection. More specifically, presented herein are techniques for presenting visualizations of an industrial facility, wherein the visualizations include one or more digital twins that respectively correspond to a machine of the industrial facility. The techniques include generating a digital twin of the machine; determining at least one property of the digital twin based on a simulation of an operation of the machine; and generating a presentation of the industrial facility that includes a visualization of the digital twin and a visual indicator of the at least one property of the digital twin.

[2862] It is to be appreciated that these techniques can be used in synergistic combinations. For example, anomaly detection and diagnosis can involve machine learning models that can automatically detect anomalies in data, wherein the machine learning models include and/or interact with one or more intelligent agent frameworks. Also, intelligent agent frameworks can involve one or more synaptic agents, and can apply the synaptic agents to monitor and/or analyze the digital twins of one or more machines in an industrial facility. Also, anomaly detection and diagnosis can involve machine learning models that can automatically detect anomalies in data, wherein the machine learning models monitor and/or analyze the digital twins of one or more machines in an industrial facility. These and other combinations of the techniques presented herein

may be synergistically combined to provide additional improvements in fields such as anomaly detection and diagnosis, intelligent agent frameworks, digital twins, and the like.

[2863] The following description provides an overview of techniques involving anomaly detection and diagnosis. It is to be appreciated that while some examples of this disclosure and the Appendices relate to machine vibration data, the disclosed techniques can also be applied to a wide range of scenarios and fields of technology, such as other mechanical anomalies (e.g., noise, heat, or electromagnetic radiation), electrical anomalies (e.g., voltage and/or current), electronic anomalies (e.g., logical errors), and the like.

[2864] In various embodiments, machine learning models can automatically detect anomalies associated with a machine (e.g., vibration data detected by vibration sensors). The techniques include recording a data set associated with the machine, determining, by a first machine learning model, a label associated with the data set, determining whether the label is to be reviewed, and responsive to determining that the label is to be reviewed, subjecting the data set and the label to a review, and updating the label based on the review.

[2865] In example embodiments, collecting the data set includes collecting data from one or more sensors associated with the machine. The sensors can include, for example and without limitation, one or more vibration sensors; one or more noise sensors; one or more temperature sensors; one or more moisture and/or humidity sensors; one or more voltage sensors; one or more current sensors; one or more electromagnetic field sensors; one or more light sensors; one or more chemical sensors; one or more pressure sensors; one or more position sensors; one or more orientation sensors; one or more data sensors; or the like. The sensors can collect data from the machine on a periodic basis (e.g., one per second, hour, or day) and/or in response to an event (e.g., upon request or in response to another event). In example embodiments, the sensors initially sample a sampled data set, which is downsampled to generate a downsampled data set that is recorded. The downsampling can include, e.g., consolidating two or more values of the data set into one value, reducing a resolution of one or more values of the data set, binning one or more values of the data set into a lower-resolution bin, applying lossy and/or lossless compression to the data set, and/or discarding one or more values of the data set that are not informative. The downsampling of the data set can reduce the volume of data to be stored and/or analyzed.

[2866] In example embodiments, the data set is used to train a machine learning model, including the first machine learning model. In example embodiments, the training occurs when a size of the data set satisfies a volume of data that is sufficient to train the machine learning model. In example embodiments, training the machine learning model includes continuing or updating a training of a previously trained and/or partially trained machine learning model (e.g., to compensate for drift or new data). In example embodiments, training the machine learning model includes training a replacement machine learning model to replace an existing machine learning model. The replacement machine learning model may include a same, similar, or different machine learning architecture as the existing machine learning model. In example embodiments, training the machine learning model includes adding the trained machine learning model to an ensemble including at least one other machine learning model.

[2867] In example embodiments, the first machine learning model is invoked synchronously to determine the label for the data set. For example, a user, such as a subject matter expert, can request an application of the first machine learning model to the data set. The system can apply the first machine learning model to the data set to determine the label for the data set, and also present the label to the user if it is to be subjected to review. The system can receive, from the user, an outcome of the review (e.g., an indication that the label is correct, an indication that the data set should have been assigned a different label associated with a different anomaly or with non-anomalous data sets, or an indication that the label is a new label associated with a new type of anomaly or non-anomalous data set). In example embodiments, the system can also present, to the user, a confidence rating of the label associated with the data set. Alternatively or additionally, in example embodiments, the system can subject the label to an automated review based on one or more review techniques. The automated review can include, for example and without limitation, applying a statistical anomaly detection analysis to the data set and at least two other data sets that are associated with the label, and characterizing the data set as an outlier as compared with the at least two other data sets that are associated with the label based on the statistical anomaly detection analysis. The automated review can include, for example and without limitation, comparing at least one statistical measurement of at least one parameter for the data set with a typical statistical measurement of the at least one parameter for the at least two other data sets. The automated review can include, for example and without limitation, applying, to the data set, a classifier model that is trained to characterize data sets as anomalies; and determining whether the label is to be reviewed based on an output of the classifier model. In these and other embodiments, based on the review of the label, the system can update the label. For example, the system can record an indication that the label is correct for the data set. The system can record an indication that a different label should have been selected for the data set instead of the label. The system can add the label to a set of existing labels in the event that the label is associated with a new type of anomaly or non-anomalous data set. In example embodiments, the system can generate and publish a report of the review of the label. In example embodiments, the system can generate, by a machine learning model, a report associated with the label, the report including one or more of, a summary of a finding associated with the label, a diagnosis of the data set associated with the label, or a next step associated with the data set. In example embodiments, the system can train a second machine learning model based on the label after updating the label based on the review. Training the second machine learning model is based on one or more of, a cluster analysis of the data set, a statistical minority oversampling of the data set, or a comparison of the label with a label selected for the data set by a subject matter expert.

[2868] In example embodiments, the first machine learning model is invoked asynchronously to determine the label for the data set. For example, the system can receive, from a user such as a subject matter expert, a label to be associated with a data set. The system can asynchronously train a machine learning model to associate the label with the data set. In example embodiments, the asynchronous training includes continuing or updating a training of a previously trained and/or partially trained machine learning model (e.g., to compensate for drift or new data). In example

embodiments, the system asynchronous training includes training a replacement machine learning model to replace an existing machine learning model. The replacement machine learning model may include a same, similar, or different machine learning architecture as the existing machine learning model. In example embodiments, the asynchronous training includes adding the trained machine learning model to an ensemble including at least one other machine learning model.

[2869] In example embodiments, the system further determines a forecast of a transition of the machine to an anomaly mode, based on the data set and the label. Alternatively or additionally, in example embodiments, the system forecasts a time point at which the machine will transition into an anomaly mode, based on the data set and the label. The forecast of the time point can be based on a linear regression model. Alternatively or additionally, in example embodiments, the system forecasts a type of anomaly mode that the machine will enter, based on the data set and the label.

[2870] Additional information about anomaly detection and diagnosis is included in Appendices of A-C of U.S. Provisional Patent App. 63/439,265, each of which is incorporated by reference as if fully rewritten herein.

Intelligent Agent Frameworks

[2871] In various embodiments, a method of deploying an agent within an environment includes associating, with the agent, at least one action, at least one instinct that motivates the agent to act within the environment, and at least one goal associated with a task within the environment, and associating, with the environment, at least one environmental rule that triggers the agent to perform the at least one action.

[2872] In example embodiments, at least a portion of the environment corresponds to and/or simulates a physical environment that exists in the physical world. In some such embodiments, the environment may include one or more parameters that correspond to one or more parameters of the physical world, such as a size parameter of a size of the physical world, an ambient parameter corresponding to an ambient feature of the physical world, or an entity that corresponds to an entity in the physical world. In example embodiments, parameters of the environment are changed (e.g., updated) to maintain correspondence with corresponding parameters of the physical world. The environment can transmit, to the agent, sensor data detected by at least one sensor within the physical world at a point in time. In example embodiments, the environment includes a mapping between one or more parameters of the physical world and one or more properties of the environment, such as a feature of an entity. The mapping can be used to maintain the correspondence between the properties of the physical world and the features of the entity within the environment.

[2873] In example embodiments, at least a portion of the environment is a simulacrum of an environment that could, but does not, physically exist in the physical world.

[2874] In example embodiments, at least a portion of the environment is a virtual environment that could not exist in the physical world. In some such embodiments, the environment corresponds to a second virtual environment, such as a digital environment. The environment may include one or more parameters that correspond to one or more parameters of the digital environment, such as a size parameter of a size of the digital environment, an ambient parameter corresponding to an

ambient feature of the digital environment, or an entity that corresponds to an entity in the digital environment. In example embodiments, parameters of the environment are changed (e.g., updated) to maintain correspondence with corresponding parameters of the digital environment. The environment can transmit, to the agent, sensor data detected by at least one sensor within the digital environment at a point in time. In example embodiments, the environment includes a mapping between one or more parameters of the digital environment and one or more properties of the environment, such as a feature of an entity. The mapping can be used to maintain the correspondence between the properties of the digital environment and the features of the entity within the environment.

[2875] In example embodiments, the one or more environmental rules include a mapping between the agent and digital environment. The one or more environmental rules can govern the functionality and possible actions that the agent can take within the environment. For example and without limitation, the environmental rules could include one or more physics rules, such as a gravity rule, a strong nuclear force rule, a weak nuclear interaction rule, an electromagnetic force rule, an inertia rule, a momentum rule, a conservation of motion rule, a conservation of energy rule, an entropy rule, and/or a causal direction of time rule. The environmental rules may include one or more physical constants, such as a gravitational constant, a speed of light constant, or a Planck constant. The environmental rules may correspond to physics laws and constants of the physical world or may diverge from the physics laws and constants of the physical world.

[2876] In example embodiments, the one or more actions associated with the agent are defined by action definitions within a mapping between the agent and the environment. The one or more action definitions can indicate, define, permit, restrict, and/or regulate the actions that the agent can perform within the environment according to the environmental rules (e.g., the principles governing movement of the agent within the physical environment, such as flight or the inertial rules of movement).

[2877] In example embodiments, the agent is further allocated an agent memory that is configured to store at least one event occurring within the environment. The agent memory may include a short-term agent memory, which stores events that occur within a short-term time frame (e.g., within the last hour or the last day). Events stored in the short-term agent memory may be discarded when a time of the event exceeds the short-term time frame. In example embodiments, the agent memory further includes a long-term agent memory that includes every event that has been stored in the short-term agent memory. In example embodiments, events stored in the long-term agent memory may be discarded based on a capacity threshold of the long-term agent memory. In example embodiments, events stored in the long-term memory may be retained, and a capacity allocated to the long-term agent memory may be expanded to accommodate additional events. In example embodiments, events in the short-term agent memory and/or the long-term agent memory may be prioritized (e.g., based on an association of the event with at least one instinct and/or at least one goal, or a clustering of events, wherein at least one distinctive event for each cluster is retained).

[2878] In example embodiments, the agent is further configured to forecast events. Such forecasts may include, for example and without limitation, a future state of at least one entity within the environment, or an outcome of an action based on at least one parameter of the environment. The forecasts can inform and/or instruct the agent as to actions to be taken. For example and without limitation, if the agent forecasts a future state of the agent that is desirable, the agent can take actions that may improve the likelihood of the forecast; and if the agent forecasts a future state of the agent that is undesirable, the agent can take actions that may reduce the likelihood of the forecast. In example embodiments, the agent is further configured to perform a determination of whether to take an action based on at least one parameter of the environment. The determination can be based on at least one of a cluster analysis of at least one label associated with the at least one parameter of the environment. The determination can be based on a classifier model that has been trained to determine a label associated with the at least one parameter of the environment.

[2879] In example embodiments, the instincts of the agent can include (for example and without limitation) an observing instinct to observe at least one property of the environment; a predicting instinct to predict at least one property of the environment; a communicating instinct to communicate with at least one other agent or entity within the environment; a prioritizing instinct to select or adjust at least one priority of the agent; a learning instinct to learn at least one feature about the environment; and/or an acting instinct to perform at least one action within the environment.

[2880] In example embodiments, the agent includes at least two goals, and is further configured to assign a priority to each goal of the at least two goals. The agent can then choose actions to be performed in pursuit of each goal and can choose each action to be performed based on the priority assigned to each goal. For example and without limitation, the agent can selectively and/or preferentially perform actions that are associated with a higher-priority goal over actions that are associated with a lower-priority goal or that are not associated with any goal. For example and without limitation, the agent can perform actions that are associated with a higher-priority goal at a greater frequency than actions that are associated with a lower-priority goal or that are not associated with any goal. The agent can update the priorities associated with the goals based on one or more properties of the environment, such as one or more properties of the agent and/or one or more properties of one or more other entities within the environment.

[2881] In example embodiments, the agent further includes a control logic that motivates the agent to perform actions associated with the at least one instinct and the at least one goal. The control logic can aid the agent in determining whether, at each point in time, to perform an action that is associated with an instinct or an action that is associated with a goal. The control logic can be based on one or more properties of the environment, such as one or more properties of the agent and/or one or more properties of one or more other entities within the environment. In example embodiments, the control logic includes a learning component that the agent can adjust and/or update based on previously performed actions and the outcomes of such actions on the instincts and goals of the agent and/or one or more properties of the environment. In example embodiments, the control logic includes a feedback loop. For example and without limitation, the agent can take

an action based on the control logic; determine an outcome of the action; based on the outcome, adjust and/or update a priority of one or more instincts and/or one or more goals; and based on the priorities, adjust and/or update the control logic to alter the selection of actions in the future.

[2882] Additional information about intelligent agent frameworks is included in Appendices of A-C of U.S. Provisional Patent App. 63/439,265, each of which is incorporated by reference as if fully rewritten herein.

Industrial Facility Digital Twins

[2883] In various embodiments, a method of presenting an analysis of a machine included in an industrial facility includes generating a digital twin of the machine; determining at least one property of the digital twin based on a simulation of an operation of the machine; and generating a presentation of the industrial facility that includes a visualization of the digital twin and a visual indicator of the at least one property of the digital twin.

[2884] In example embodiments, the industrial facility includes one or more of a resource mining facility, a resource measurement and/or management facility, a manufacturing plant, a chemical processing facility, an energy generation facility, a waste processing facility, a storage facility, or the like. In example embodiments, the industrial facility is scanned to generate a digital representation of the industrial facility. For example and without limitation, the scan may be captured by a still camera, a video camera, a LIDAR sensor, or the like. In example embodiments, the scan is based on at least one sensor input received by at least one environmental sensor within the industrial facility. In example embodiments, the machine includes one or more of a lever, a pulley, a flywheel, a gasoline-powered vehicle, a motor, a rotor, a stator, a fan, an electric vehicle, a server, a workstation, a piece of communication equipment, or the like.

[2885] In example embodiments, the digital twin includes a digital representation of the machine that can be digitally manipulated to simulate one or more operations of the machine. For example and without limitation, the machine can include an arrangement of one or more mechanical, chemical, electrical, and/or electronic components, and the digital twin can include a digital representation of the arrangement of digital representations of the corresponding components. In example embodiments, the digital twin is generated from a schematic depiction of the machine. In example embodiments, the digital twin is generated from a scan of the machine. For example and without limitation, the scan may be captured by a still camera, a video camera, a LIDAR sensor, or the like. In example embodiments, the digital twin is based on at least one sensor input received by at least one sensor associated with the machine within the industrial facility.

[2886] In example embodiments, the digital twin of the machine is retrieved from a digital twin library. In example embodiments, a machine type of the machine is determined, and a stored representation of digital twins of machines of the machine type is retrieved from the digital twin library. In some cases, one or more features of the machine in the industrial facility differs from corresponding one or more features of the retrieved digital twin. As a first such example, various instances of a particular machine type can vary due to variations in the manufacturing materials and/or processes, resulting in machines of a same machine type but with different features. As a first such example, the digital twin library might not include a stored representation of machines

of the same machine type as the machine in the industrial facility, but might include a stored representation of a similar machine type (e.g., an earlier, later, or alternative version of the machine, or a similar machine from another vendor). In such cases, the at least one feature of the machine can be determined, and at least one corresponding feature of the retrieved digital twin can be adjusted based on the at least one feature of the machine.

[2887] In example embodiments, various operations of the machine can be simulated. For example and without limitation, the simulated operations can include designed and/or archetypal operations of the machine according to the machine design specifications, actual operations that correspond to actual operations of the machine in the industrial facility (which may differ from the designed and/or archetypal operations of the machine), and/or anomalous operations of the machine due to an anomaly state. The operation can include one or more machine tasks such as calibration, a manufacturing operation, a physical and/or chemical processing operation, an electrical power charging, discharging, and/or exchange operation, a cleaning operation, a self-maintenance operation, a diagnostic operation, an idle operation during an idle state, or the like. In example embodiments, the simulation of the operation of the machine is based on a vibration generated by at least one moving part of the machine and detected by a vibration sensor.

[2888] In example embodiments, the simulation may include the simulation of at least two components of the machine. For example and without limitation, the simulation can include a simulation of an interaction between a motion of a first component of the machine and a second (moving or non-moving) component of the machine.

[2889] In example embodiments, the simulation includes the simulation of a second machine of the industrial facility. The simulation can include a simulation of an interaction between the machine and the second machine of the industrial facility. For example and without limitation, vibration generated by a first machine of the industrial facility might affect an operation of a second machine of the industrial facility, and the simulation can include a simulation of the effect of the vibration generated by the first machine on an operation of the second machine.

[2890] In example embodiments, the presentation of the industrial facility is generated by rendering, by a graphics engine that renders three-dimensional digital environments, a three-dimensional digital environment that corresponds to the industrial facility, and the visualization of the digital twin is based on a three-dimensional digital model of the machine rendered by the graphics engine. For example and without limitation, the graphics engine can include a game engine, such as the UNITY® game engine or the UNREAL ENGINE® game engine. Many such graphics engines include a physics engine that also simulates physical interactions between three-dimensional models, such as collisions, forces, surface deformation, light reflection and/or refraction, and the like. In some such embodiments, the simulation of the operation of the machine is performed by the physics engine included in the graphics engine.

[2891] In example embodiments, the visualization of the digital twin can include one or more types of visualizations. For example and without limitation, the visualization can include at least one of, a visual indicator of a location of the at least one property, a visual indicator of a measurement of the at least one property, a visual indicator of a vibration associated with the at least one property,

a text description of the at least one property, a visual indicator of a classification of the at least one property, a media recording of a corresponding property of the machine, a label associated with the at least one property, a visual depiction of the operation of the machine in the industrial environment, or a visual depiction of the simulation of the operation of the machine.

[2892] Additional information about industrial facility digital twins is included in Appendices of A-C of U.S. Provisional Patent App. 63/439,265, each of which is incorporated by reference as if fully rewritten herein.

Programmable Logic Controllers

[2893] In example embodiments, an infrastructure includes one or more programmable logic controllers (PLCs) that apply programmable logic to the data, processes, devices, and/or resources of the infrastructure. By way of example, infrastructures that may include one or more PLCs include factories and other industrial environments; office environments; residential environments, such as housing complexes or neighborhoods; agricultural environments; governmental and/or municipal environments; energy grids; vehicles and vehicle networks, including (for example) trains, automobiles, airplanes, and space-faring vehicles such as satellites; and the like, including combinations thereof. In such environments, PLCs may perform various logical operations, including (for example) receiving data, detecting data by various sensors (e.g., cameras, microphones, tactile input devices, voltage sensors, or the like), receiving and/or transmitting data by a wired and/or wireless network, receiving and/or transmitting data by a device-to-device communication channel (e.g., infrared or BLUETOOTH™), filtering data, analyzing data, generating data, requesting data, retrieving data, executing code, storing data, reporting data, displaying data, or the like.

[2894] In example embodiments, one or more PLCs may physically exist and execute in proximity to an infrastructure. For example, a factory may include a room including industrial equipment that is managed by one or more physically collocated PLCs. Alternatively or additionally, one or more PLCs may physically exist and execute remotely from an infrastructure. For example, a factory may include a room including industrial equipment that is managed by one or more PLCs that are located apart from a factory, such as adjacent to the factory or distant from the factory. Alternatively or additionally, one or more PLCs that manage an infrastructure may be virtualized and/or simulated. For example, one or more PLCs may be embodied as a set of instructions that are executed in a virtual environment that is executed by hardware that is the same as, similar to, or different than the virtualized hardware of the one or more PLCs.

[2895] In example embodiments, one or more PLCs may be included in a local area network (LAN) that includes at least a portion of the infrastructure, wherein the one or more PLCs and the at least of the infrastructure are physically collocated and/or physically interconnected. Alternatively or additionally, one or more PLCs may be included in a virtual network, such as a virtual private network (VPN) that logically groups a set of resources that may or may not be physically collocated and/or physically interconnected. For example, a virtual private cloud (VPC) may include one or more devices, machines, resources, or the like of an infrastructure and one or more PLCs that manage the one or more devices, machines, resources, or the like of the

infrastructure, wherein such one or more PLCs may be physically collocated and/or physically interconnected with the infrastructure and/or may be physically distant from and/or only remotely interconnected with the infrastructure.

[2896] In example embodiments, one or more PLCs may be programmed and/or configured by a programmer. For example, one or more programmers may develop instructions in a programming language. Such programming languages may include (for example) one or more industrial programming languages, such as Structured Text (ST), Sequential Function Charts (SFC), Ladder Logic Diagrams (LD), Function Block Diagrams (FBD), Instruction List (IL) or the like. Alternatively or additionally, such programming languages may include (for example) one or more general-purpose programming languages, such as C, C++, C#, Java, Python, or the like. Alternatively or additionally, such programming languages may include (for example) one or more device-specific, hardware-descriptive, and/or low-level programming languages, such as assembly, the VHSIC Hardware Description Language (VHDL), or VERILOG. Alternatively or additionally, one or more programmers may develop the instructions in a non-programming language, including a natural language (e.g., English or Spanish) and/or a synthetic language (e.g., Esperanto). Alternatively or additionally, one or more programmers may provide instructions in a first language (e.g., Python) that are then translated, transpiled, or the like into a second language (e.g., C) for further processing (e.g., further translation, compilation, interpretation, *etc.*) by the PLC.

[2897] In example embodiments, one or more programmers may develop the instructions on a PLC for execution by a PLC or a collection of PLCs. Alternatively or additionally, one or more programmers may develop the instructions on a first device (e.g., a first PLC or a collection of first PLCs) for execution by a second PLC or a collection of second PLCs. Alternatively or additionally, one or more programmers may provide first instructions on a first device (e.g., a first PLC or a collection of first PLCs), and, based on the first instructions, the first device may generate second instructions for execution by a second PLC or a collection of second PLCs.

[2898] In example embodiments, the developer may develop the instructions in a familiar programming interface, such as a text editor or an integrated development environment ("IDE"). Alternatively or additionally, the developer may develop the instructions in a "low-code" or "no-code" programming interface, such as a visual drag-and-drop interface in which the logic of the instructions is arranged as a series of interconnected visual blocks. Alternatively or additionally, the developer may develop the instructions via a natural language interface (e.g., by specifying verbal commands to a device in a natural language, such as English, that the device processes to determine the instructions). In example embodiments, instructions may be developed by a plurality of developers working together (e.g., as a team) and/or independently (e.g., a synthesis of instructions by a first developer and a second developer). Alternatively or additionally, the instructions may be developed wholly or partially by a device or a component thereof (e.g., a generative machine learning model), with or without the input of a developer.

[2899] In example embodiments, instructions may be provided to a PLC in wholly or partially compiled form, such as bytecode targeting a particular architecture of the PLC or intermediate

language (IL) code. Alternatively or additionally, the instructions may be provided to the PLC in wholly or partially uncompiled form, and may be wholly or partially compiled by the PLC for execution. Alternatively or additionally, the instructions may be provided to the PLC in uncompiled form, and may be executed by the PLC in uncompiled form (e.g., as an interpreted script).

[2900] In example embodiments, a PLC may execute received and/or generated instructions and then dispose of the instructions. For example, a PLC may receive a stream of instructions from another device (e.g., another PLC), may execute the instructions upon receipt, and may dispose of the instructions after execution. The instructions may be provided to the PLC on a just-in-time basis (e.g., in response to an event, and/or for immediate or at least prompt execution by the PLC). Alternatively or additionally, a PLC may store a received and/or generated set of instructions for execution in the future. For example, the instructions may be provided to the PLC for scheduled and/or triggered execution (e.g., at a specific date, time, and/or interval, and/or in response to a detection of an event in the future). In example embodiments, the instructions may be provided to the PLC for conditional execution (e.g., in the event that a particular set of criteria are satisfied, and/or on the volition of the PLC).

[2901] In example embodiments, a PLC may interoperate with one or more other devices, processes, or other resources. Such interoperation may occur before, during, and/or after executing the instructions. For example, a PLC may communicate with a second device (e.g., a second PLC or a collection of other PLCs) to execute the instruction set. At least a portion of the instruction set may be partially executed by the PLC and also by another device, including another PLC, either independently or cooperatively. Alternatively or additionally, a first PLC may execute a first portion of the instruction set, and a second device (e.g., a second PLC) may execute a second portion of the instruction set before, during, and/or after the execution of the first portion of the instruction set by the first PLC. The first PLC may execute a portion of the instruction set in response to an execution of a portion of the instruction set by a second device (e.g., a second PLC), or *vice versa*. A PLC may execute a first portion of the instruction set concurrently with an execution of a second portion of the instruction set by the same PLC and/or another device (e.g., simultaneously and/or interleaved). A PLC may execute a first portion of the instruction set in response to an execution of a second portion of the instruction set by the same PLC and/or another device (e.g., as a condition and/or result of executing the second portion of the instruction set).

[2902] In example embodiments, one or more PLCs include one or more forms of artificial intelligence (AI). The AI may include one or more machine learning models, such as an artificial neural network (ANN), decision tree, expert system, or the like. The PLC may include the AI in fixed hardware (e.g., an AI-based ASIC), programmable hardware (e.g., an FPGA), firmware (e.g., a flash memory including the hyperparameters and parameters of an artificial neural network), software (e.g., volatile memory storing the hyperparameters and parameters of an artificial neural network), a non-transitory storage medium (e.g., a hard disk drive storing the hyperparameters and parameters of an artificial neural network), or the like. The PLC may include one or more digital twins that correspond to one or more physical devices, machines, resources, entities, individuals,

processes, or the like. The PLC may include one or more intelligent agents that create, use, interact with, manage, and/or transport one or more physical devices, machines, resources, entities, individuals, processes, or the like. Such forms of artificial intelligence may be wholly or partially stored and/or executed by the PLC (e.g., an artificial neural network may be distributed over and cooperatively executed by the PLC and a second device, such as a second PLC). Alternatively or additionally, the PLC may communicate with another device that stores and/or executes at least a portion of a form of AI (e.g., a second PLC that stores and/or executes at least a portion of an artificial neural network). The PLC, including a form of AI stored and/or executed thereby, may receive data and/or instructions from one or more other devices or components thereof, such as a digital twin or agent executed by a second PLC.

[2903] In example embodiments, one or more PLCs may receive data from one or more sources and may filter the data to a subset that is relevant to one or more processes, such as detection, analysis, summarization, reporting, curation, presentation, alerting, or the like. Such sources may include (for example) public data repositories, such as public webservers or public databases; private data repositories, such as private webservers or private databases; internal data repositories, such as webservers or databases of the industrial infrastructure managed by the one or more PLCs; other components of the one or more PLCs (e.g., output of a machine learning model or digital twin executing on the one or more PLCs); and/or devices or other components thereof (e.g., output of a machine learning model or digital twin executing on one or more other PLCs). Alternatively or additionally, one or more PLCs may generate data (e.g., by a generated AI model).

[2904] In example embodiments, one or more PLCs may process received and/or generated data by one or more components (e.g., a machine learning model or digital twin executing on the one or more PLCs). The one or more PLCs may generate one or more outputs of the processing, such as executing additional instructions in response to the processing; using an output of the processing as an input to another process; storing an output of the processing; reprocessing, iteratively processing, and/or continuously processing an output of the processing (e.g., by a reentrant and/or recurrent neural network (RNN)); transmitting an output of the processing to another device; and/or generating an output or presentation based on the processing (e.g., generating a visualization or alert based on an output of the processing). In example embodiments, the processing of the PLC manages one or more resources of an infrastructure, such as (for example) the PLC itself, one or more other PLCs, machines, robots, individuals, entities, or the like. For example, the processing of the PLC may generate, schedule, start, stop, configure, adjust, control, monitor, report on, or otherwise affect an industrial process of an industrial infrastructure. Alternatively or additionally, the processing of the PLC may cause a programming or reprogramming of the PLC itself or one or more other PLCs.

Generative Artificial Intelligence for Industrial Internet of Things

[2905] In example embodiments, a generative artificial intelligence engine (GAIE) may be combined with a machine learning system in a transaction environment. Input to the GAIE may include images, video, audio, text, programmatic code, data, and the like. Outputs from a GAIE may include structured and organized prose, images, video, audio content, software / programming

source code, formatted data (e.g., arrays), algorithms, definitions, context-specific structures (e.g., smart contacts, transaction platform configuration data sets, and the like), machine language-based data (e.g., API-formatted content), and the like. For GAIE instances in which the models are designed to process text data, the GAIE may interface to other programmatic systems (such as traditional machine learning engines) to process other forms of data into text data. In example embodiments, the other programmatic systems, including systems executing machine learning algorithms, may produce textual based (optionally at volume) that may be consumed by GAIE. For example, consider such another system building a series of one thousand text-based observations on the other-formatted data; this may be a useful input for a GAIE model to learn and process (e.g., summarize) into text-formatted output information. In example embodiments, an interface between the GAIE and its combined machine learning system may be extended to include a dialogue between the systems, where the GAIE includes and/or accesses a capability to ask the machine learning system specific questions to facilitate the refining of its knowledge. For example, the dialogue capability may include a request of the machine learning system to provide an assessment of current market trading positions. In another example, the dialogue capability may encode numeric outputs from the machine learning engine into text (e.g., words, such as high, medium, low) that may be input for interpretation by the GAIE.

[2906] In example embodiments, the data processed by a GAIE may include one or more types of content. For example, a GAIE may receive, as input, data that represents one or more natural-language expressions, single- or multidimensional shapes or models, real-world and/or virtual scene representations, LIDAR point-cloud representations, sensor inputs and/or outputs, vehicle and/or machine telemetry, geographic maps, authentication credentials, financial transactions, smart contracts, processing directives and/or resources such as shaders, device configurations such as HDL specifications for programming FPGAs, databases and/or database structural definitions, or the like, including metadata associated with any such data types. Input to the GAIE may also include data that represents one or more features of another machine learning model, such as a configuration (e.g., model type, parameters, and/or hyperparameters), input, internal state (e.g., weights and biases of at least a portion of the model), and/or output of the other machine learning model. These and other forms of content may be received as various forms of data. For example, a natural-language expression received as input by a GAIE could be encoded as one or more of: encoded text, an image of a writing, a sound recording of human speech, a video of an individual exhibiting sign language, an encoding according to a machine learning model embedding, or the like, or any combination thereof. In example embodiments, an input received and processed by the GAIE can include an internal state of the GAIE, such as a partial result of a partial processing of an input, or a set of weights and/or biases of the GAIE as a result of prior processing (e.g., an internal state of a recurrent neural network (RNN)).

[2907] In example embodiments, the data and/or content received and processed by a GAIE originates from one or more individuals, such as a person speaking a natural-language expression. In example embodiments, the data and/or content received and processed by a GAIE originates from one or more natural sources, such as patterns formed by nature. In example embodiments, the

data and/or content received and processed by a GAIE originates from one or more other devices, such as another machine learning model executing on another device, or from another component of the same device executing the GAIE, such as output of another machine learning model executing on the same device executing the GAIE, or a sensor in an Internet of Things (IoT) and/or cloud architecture. In example embodiments, the data and/or content received and processed by a GAIE is artificially synthesized, such as synthetic data generated by an algorithm to augment a training data set. In example embodiments, the data and/or content received and processed by a GAIE is generated by the same GAIE, such as an internal state of the GAIE in response to previous and/or concurrent processing, or a previous output of the GAIE in the manner of a recurrent neural network (RNN).

[2908] In example embodiments, at least some or part of the data and/or content received and processed by a GAIE is also used to train the GAIE. For example, a variational GAIE could be trained on an input and a corresponding acceptable output, and could later receive the same input in order to output one or more variations of the acceptable output. In example embodiments, at least some or part of the data and/or content received and processed by a GAIE is different than data and/or content that was used to train the GAIE. In some such embodiments, the data and/or content received and processed by the GAIE is different than but similar to the data and/or content that was used to train the GAIE, such as new inputs that are exhibit a similar statistical distribution of features as the training data. In some such embodiments, the data and/or content received and processed by the GAIE is different than and dissimilar to the data and/or content that was used to train the GAIE, such as new inputs that exhibit a significantly different statistical distribution of features than the training data. In scenarios that involve dissimilar inputs, one or more first outputs of the GAIE in response to a new input may be compared to one or more second outputs of the GAIE in response to inputs of the training data set to determine whether the first outputs and the second outputs are consistent. The GAIE may request and/or receive additional training based on the new inputs and corresponding acceptable outputs. In scenarios that involve dissimilar inputs, the GAIE may present an alert and/or description that indicates how the new inputs and/or corresponding outputs differ from previously received inputs and/or corresponding outputs.

[2909] In example embodiments, the output of a GAIE may include one or more types of content. For example, a GAIE may generate, as output, data that represents one or more natural-language expressions, single- or multidimensional shapes or models, real-world and/or virtual scene representations, LIDAR point-cloud representations, sensor inputs and/or outputs, vehicle and/or machine telemetry, geographic maps, authentication credentials, financial transactions, smart contracts, processing directives and/or resources such as shaders, device configurations such as HDL specifications for programming FPGAs, databases and/or database structural definitions, or the like, including metadata associated with any such data types. Output of the GAIE may also include data that represents one or more features of another machine learning model, such as a configuration (e.g., model type, parameters, and/or hyperparameters), input, internal state (e.g., weights and biases of at least a portion of the model), and/or output of the other machine learning model. These and other forms of content may be generated by the GAIE as various forms of data.

For example, a natural-language expression generated as output by the GAIE could be encoded as one or more of: encoded text, an image of a writing, a sound recording of human speech, a video of an individual exhibiting sign language, an encoding according to a machine learning model embedding, or the like, or any combination thereof. In example embodiments, an output of the GAIE can include an internal state of the GAIE, such as a partial result of a partial processing of an input, or a set of weights and/or biases of the GAIE as a result of prior processing (e.g., an internal state of a recurrent neural network (RNN)).

[2910] In example embodiments, a language-based dialogue-enabled GAIE may be configured to produce (e.g., write) new machine learning models that may process various types of data to provide new and extended text input for processing by the GAIE. In example embodiments, humans may observe and interact with this ongoing dialogue between the two systems. In example embodiments, the dialogue is initiated by an expression of a conversation partner (e.g., a human or another device), and the GAIE generates one or more expressions that are responsive to the expression of the conversation partner. In example embodiments, the GAIE generates an expression to initiate the dialogue, and further responds to one or more expressions of the conversation partner in response to the initiating expression. In example embodiments, the ongoing dialogue occurs in a turn-taking manner, wherein each of the conversational partner and the GAIE generating an expression based on a previous expression of the other of the conversation partner and the GAIE. In example embodiments, the ongoing dialogue occurs extemporaneously, with each of the conversation partner and the GAIE generating expressions irrespective of a timing and/or sequential ordering of previous and/or concurrent expressions of the conversation partner and/or the GAIE.

[2911] In example embodiments, the dialogue occurs between a GAIE and a plurality of conversation partners, such as two or more humans, two or more other GAIEs, or a combination of one or more humans and one or more other GAIEs. In some such example embodiments, the GAIE and each of the other conversation partners take turns generating expressions in response to prior expressions from the GAIE and the other conversation partners. In some such embodiments, one or more sub-conversations occur among one or more subsets of the GAIE and the plurality of conversation partners. Such sub-conversations may occur concurrently (e.g., the GAIE concurrently engages in a first conversation with a first conversation partner and a second conversation with a second conversation partner) and/or consecutively (e.g., the GAIE concurrently engages in a first conversation with a first conversation partner, followed by a second conversation with a second conversation partner). Such sub-conversations may involve the same or similar topics or expressions (e.g., the GAIE may present the same or similar conversation-initiating expression to each of a plurality of conversation partners, and may concurrently engage each of the plurality of conversation partners in a separate conversation on the same or similar topic). Such sub-conversations may involve different topics or expressions (e.g., the GAIE may present different conversation-initiating expressions to each of a plurality of conversation partners, and may concurrently engage each of the plurality of conversation partners in a separate conversation on different topics). In example embodiments, a first conversation among a first

subset of the GAIE and conversation partners may be related to a second conversation among a second subset of the GAIE and conversation partners (e.g., the second subset may engage in a second conversation based on content of the first conversation among a first subgroup).

[2912] In example embodiments, one or more of the GAIE and the conversation partner may embody one or more roles. For example, the GAIE may generate expressions based on a role of a conversation starter, a conversation responder, a teacher, a student, a supervisor, a peer, a subordinate, a team member, an independent observer, a researcher, a particular character in a story, an advisor, a caregiver, a therapist, an ally or enabler of a conversation partner, or a competitor or opponent of a conversation partner (e.g., a "devil's advocate" that presents opposing and/or alternative viewpoints to a belief or argument of a conversation partner). In example embodiments, at least one of the one or more conversation partner embodies one or more aforementioned roles or other rules. In example embodiments, a role of a GAIE is relative to a role of a conversation partner (e.g., the GAIE may embody a superior, peer, or subordinate role with respect to a role of a conversation partner). In example embodiments, a role of a GAIE in a first conversation among a first subset of the GAIE and a plurality of conversation partners may be the same as or similar to a role of a GAIE in a second conversation among a first subset of the GAIE and the plurality of conversation partners. In example embodiments, a role of a GAIE in a first conversation among a first subset of the GAIE and a plurality of conversation partners may differ from a role of a GAIE in a second conversation among a first subset of the GAIE and the plurality of conversation partners (e.g., the GAIE may embody a role of a teacher in a first conversation and a role of a student in a second conversation). In example embodiments, a role of a GAIE in a conversation may change over time (e.g., the GAIE may first embody a role of a student in a conversation, and may later change to a role of a teacher in the same conversation). In example embodiments, a GAIE may embody two or more roles in a conversation (e.g., the GAIE may exhibit two personalities in a conversation that respectively represent one of two characters in a story). In example embodiments, a GAIE generates expressions between two or more roles in a conversation (e.g., the GAIE may generate a dialogue between each of two characters in a story). In example embodiments, a GAIE may engage in each of multiple conversations in a same or similar modality (e.g., engaging in multiple text-based conversations concurrently). In example embodiments, a GAIE may engage in each of multiple conversations in different modalities (e.g., engaging in a first conversation via text and a second conversation via voice).

[2913] In example embodiments, a GAIE participating in a conversation is associated with an avatar (e.g., a name, color, image, two- or three-dimensional model, voice, or the like). Expressions generated by the GAIE may be presented as if originating from the GAIE (e.g., in the voice associated with the GAIE, or in a speech bubble that is displayed near a visual position of a GAIE in a virtual or augmented-reality environment). In example embodiments, an avatar of a GAIE may be based on a role of the GAIE (e.g., a GAIE embodying a role of a teacher may be associated with an avatar depicting a teacher). In example embodiments, an avatar of a GAIE may be included in a real-world actor, such as a robot in a real-world environment such as a stage performance.

[2914] In example embodiments, a GAIE may include generative pretrained transformer elements that may be configured as a language model designed to understand various types of input and produce chat commands for a chat-type interface system. These commands may include software development tasks, API calls, and the like. In example embodiments, such a language model may include input functions that support receiving images, including video, to build textual output, functions, and additional questions that may be injected into the dialogue between the two systems in the dialogue embodiment described above. In example embodiments, this multimodal support may allow for contextual analysis of images and other media formats. In an example, users/customers may upload images or other media into a GAIE enabled platform. Based on aspects of a corresponding input prompt, a multi-modal GAIE may be configured for use in a valuation workflow to identify both macro and micro attributes and their correlated effects on valuation from a plurality of perspectives. In this example, photographs/images of an old car may be input along with a valuation-related prompt. In response, the GAIE may identify one or more typical values based on detected attributes of the car, such as the make/model, etc. The GAIE may further take into account finer details in the image to suggest potential value-altering metrics. In one example, a finer detail in the image such as damaged body panels may reduce the car value below a typical value. In another example, a finer detail in the image that shows a marking consistent with a limited production run may increase the valuation.

[2915] In example embodiments, such a GAIE may facilitate workflow orchestration for a process that uses a conversational, generative AI agent and another AI-supported process in an orchestrated sequence. In example embodiments, a GAIE may generate, perform, maintain, and/or supervise one or more workflows in a robotic process automation (RPA) environment. For example, a GAIE may be trained to monitor expressions and/or actions of an individual during interaction with other individuals, and may generate similar expressions and/or perform similar actions during similar interactions between the GAIE and other individuals. In some such scenarios, the GAIE passively observes the individual during the interactions with other individuals and self-trains to behave similarly to the individual in similar interactions with other individuals. In some such scenarios, the individual actively trains and/or teaches the GAIE to generate expressions and/or actions (e.g., by creating and/or performing example or pedagogical interactions the GAIE), and based on the training and/or teaching, the GAIE behaves similarly during subsequent interactions between the GAIE and other individuals. In example embodiments, the GAIE is trained and/or taught by an individual to perform a behavior while interacting with individuals, and subsequently performs the behavior while interacting with the same individual who provided the training and/or teaching.

[2916] In example embodiments, a pretrained GAIE system may have a smart contract analysis engine that determines one or more features of a smart contract that is under consideration by a user. The GAIE may further have a conversation engine that explains the features of the smart contract to the user, including summarizing contents of smart contracts.

[2917] In example embodiments, a GAIE may be pre-trained to perform prompt generation based on a data story or a plurality of sources across systems. Example generated prompts may include instructing and/or requesting the pre-trained GAIE to tell a story about a journey of a product, a

business relationship, an event, a service provider, a smart container fleet, a robotic fleet, and the like.

[2918] In example embodiments, the GAIE may receive a plot or outcome of the story, and may generate content that is content with the plot or that produces the outcome. In example embodiments, the GAIE may generate a plot or outcome of the story, and may also generate content that is consistent with the GAIE-generated plot or outcome of the story. In example embodiments, the GAIE may receive a world or environment of a story, and may generate content that occurs within the given world or environment. In example embodiments, the GAIE may generate a world or environment of a story, and may also generate content that occurs within the GAIE-generated world or environment. In example embodiments, the GAIE may receive a character or event to be included in a story, and may generate content that includes the given character or event in the story. In example embodiments, the GAIE may generate a character or event to be included in a story, and may also generate content that includes the GAIE-generated character or event in the story. In example embodiments, the GAIE may generate a world, environment, character, event, or the like "from scratch" (e.g., based on randomized inputs). In example embodiments, the GAIE may generate a world, environment, character, event, or the like based on a given world, environment, character, event, or the like (e.g., a story that is based on a real-world public figure or event).

[2919] In example embodiments, the GAIE may receive a first story and may generate a second story that is related to the first story. For example, the GAIE may generate a second story that is an alternative retelling of the first story (e.g., a second story that includes a retelling of the first story from a perspective of a different character than a narrating character of the first story). The GAIE may generate a second story that occurs in a same or similar world or environment as the first story, or a different world or environment that is related to a world or environment of the first story. The GAIE may generate a second story that features a character or event of the first story, or a different character or event that is related to a character or event of the first story.

[2920] In example embodiments, the GAIE may generate a story from the perspective of a narrator or independent observer of the story (e.g., a third-person story). In example embodiments, the GAIE may generate a story from the perspective of a character or point of view within the story (e.g., a first-person story), including a character generated and/or embodied by the GAIE. In example embodiments, the GAIE may generate a story from the perspective of a listener or audience member to whom the story is presented (e.g., a second-person story). In example embodiments, the GAIE may generate a story from multiple perspectives, such as a first part of a story generated from a perspective of a first character, a second part of the story generated from a perspective of a second character, and a third part of a story generated from a perspective of a narrator. In example embodiments, the GAIE may generate a story involving a sequence of two or more events (e.g., a story that involves two or more events observed by a character). In example embodiments, the GAIE may generate a story involving an event that is portrayed from multiple perspectives (e.g., a story that describes an event from a perspective of a first character, and that also describes the same event from a perspective of a second character).

[2921] In example embodiments, a GAIE may generate a static story that remains the same upon retelling. In example embodiments, the GAIE may generate a dynamic story that changes upon retelling (e.g., adding more detail to a story upon each retelling). In example embodiments, a GAIE may change a story based on an input of a user (e.g., based on a choice of outcomes selected by one or more receivers of the story). In example embodiments, a GAIE may generate a story based on one or more inputs received from one or more receivers of the story (e.g., based on a prompt of a user, such as a request to create a story that includes a certain event specified by the user). In example embodiments, a GAIE may receive feedback from a receiver about a story (e.g., an expression of pleasure, displeasure, approval, disapproval, delight, dissatisfaction, confusion, or the like regarding a character, event, or property of the story), and the GAIE may update the story based on the feedback (e.g., adding, removing, or clarifying an event in the story, or switching a perspective of an event from a first character in the story to a second character in the story).

[2922] In example embodiments, a GAIE may be pre-trained for operation in an industrial Internet-of-Things (IoT) environment, such as facilitating computer vision. In an example, images processed by a GAIE-enabled system may cause the GAIE to identify potential defects in a product. In another example, images processed by a GAIE-enables system may analyze a scene (e.g., an area of a road or road network that includes a number of vehicles) and may generate a description of the analysis of the scene (e.g., a description of what is occurring in the scene, such as a cause of traffic or a vehicular accident). In another example, images processed by a GAIE-enables system may analyze a scene (e.g., an area of a road or road network that includes a number of vehicles) and may generate an output associated with the scene (e.g., a control signal for traffic control device within the area based on a traffic pattern of the area, a control signal for a vehicle operating within the area, or a routing signal for a vehicle that is or may be traveling through or near the area). Factors that may impact pretraining may include accuracy, consistency, granularity, latency, cost, and a requirement for proprietary training data (e.g., a volume of training data and/or a set of labels to be associated with the training data). Whereas a computer vision application might apply fast edge processing (that outpaces what a generative AI contextualizer may provide), a pre-trained GAIE may apply next-token generation processing, such as described herein, to determine aspects of an image that fit to a set of pre-trained prompts/responses regarding potential defects. For example, a GAIE can receive, as input, one or more images associated with a scene and one or both of (i) a previous output of the GAIE associated with the scene (e.g., a control input previously provided to a vehicle that contributed to the scene) or (ii) a current output of the same GAIE or another machine learning model (e.g., a description of the same or similar scene as previously or concurrently generated by the GAIE or another machine learning model). The GAIE can analyze the previous and/or current output related to the scene to determine an accuracy, validity, or quality of the output (e.g., as an adversarial discriminator of the GAIE or another machine learning model).

[2923] In example embodiments, a GAIE may be trained by loading data (such as structured and un-structured data that may be dominated by numerical or non-text values) to the GAIE. Examples of such training data may include one or more database schemas. Techniques for curation and

integration of purpose-specific data, including curation of models as inputs to a GAIE may include curating domain-specific data, data and model discovery.

[2924] Candidate areas of innovation enabled by and/or associated with GAIE advances may include user behavior models (optionally with feedback and personalization), group clustering and similarity, personality typing, governance of inputs and process, explaining the basis of GAIE knowledge and proof points, genetic programming with feedback functions, intelligent agents, voice assistants and other user experiences, transactional agents (counterparty discovery and negotiation), agents that deal with other agents, opportunity miners, automated discovery of opportunities for agent generation and application, user interfaces that adapt to the user and context, hybrid content generation, collaboration units of humans and generative AI, purpose-specific data integration, a selected set of data sources, curation of data as models as input to generative AI, and the like.

[2925] In embodiments of a GAIE-enabled system, such as one for robotic process automation, the GAIE system may summarize a set of actions being subjected to robotic automation and describe context for the actions, such as, "I found these properties as fitting your criteria because of the following features. Which ones are most attractive?" In this way, a process automation system enabled with GAIE may solicit feedback for faster feedback-based training.

[2926] In example embodiments, emerging capabilities of GAIE technology may greatly improve upon earlier versions in terms of, for example, integration of domain-specific knowledge (e.g., math) with a chat interface. Further emerging capabilities may include being better informed about and for processing prompts of complex topics. Yet further, knowledge organization is becoming much improved as GAIE systems evolve. In example embodiments, updated GAIEs may correctly answer a prompt asking about today's date, whereas prior versions may answer that today's date (e.g., the current date) may be the date on which the GAIE was last trained.

[2927] In example embodiments, a context pretrained (e.g., subject matter focused) GAIE may provide better personalization than a base GAIE instance. In general, while a base GAIE, if explicitly informed of details of the user may attempt to personalize its responses, a subject matter focused or other pre-trained GAIE may be configured with and/or with access to structured information about users (e.g., determined based on user identification and/or prompt-based clues, and the like) to provide inherent, latent context for a dialogue that includes user personalized responses.

[2928] In example embodiments, a GAIE is configured to support interpretability and/or explainability of its outputs. In example embodiments, a GAIE provides, along with an output, a description of a basis of the output, such as an explanation of the reason for generating this particular output in response to an input. In example embodiments, a GAIE provides, along with an output, a description of an internal state of the GAIE that resulted in the output, such as a set of variational parameters of a variational encoder that were processed in combination with an input to produce an output, and/or an internal state of the GAIE due to a previous processing of the GAIE that resulted in the output (e.g., similar to a recurrent neural network (RNN)). In example embodiments, a GAIE provides, along with an output, an indication of one or more subsets of

features of an input that are particularly associated with the output (e.g., in a GAIE that outputs a caption or summary of an image, the GAIE can also identify the particular portions or elements of the image that are associated with the caption or portions of the summary).

[2929] In example embodiments, an advanced GAIE, such as one pretrained for subject matter specific operation, may be trained for improved epistemology, to help determine evidence of the content that it represents as facts in responses that it provides. One example of improved epistemology may include citing sources of knowledge pertinent to facts in a response as a step toward proof of facts of a response - essentially a way of the GAIE "showing its work," or at least where its work originates. In example embodiments, a GAIE generates output based on information received from one or more external sources (e.g., one or more messages in a message set, or one or more websites on the Internet), and the GAIE indicates one or more portions of the information that are associated with the output (e.g., one or more websites on the Internet that provided information that is included in the output of the GAIE).

[2930] An advanced GAIE as described and envisioned herein may maintain contextual awareness across chat (user-prompt/GAIE-response) interactions. Maintaining contextual awareness may help avoid the GAIE beginning each chat session from scratch, with no context as to prior chats with the same user. Maintaining contextual awareness may also enable picking up and resuming a conversation from earlier interactions between the GAIE and a user. Yet further maintaining contextual awareness and awareness of passage of time between interaction sessions may facilitate adapting responses to prompts in a later resumed chat session based on trained knowledge of the intervening passage of time and/or changing circumstances. In an example, a GAIE may determine that a deadline described in an earlier chat has expired, a consequential intervening event has occurred (your home-town team lost the big game), and the like. Further, contextual awareness across time-separate chat sessions may be highly valuable when being employed for projects that may have real-world physical constraints on time (e.g., smart contract negotiation may involve human evaluation, discussion, and decision making that may take time based, for example on other priorities seeking involvement of the human). This may determine the difference between treating each conversation as individual/compartmentalized/isolated, and treating ongoing, time-separated conversations as resumable, optionally as if (almost) no time had passed. In example embodiments, a GAIE may be configured with a contextualization module that maintains some notion of conversation sessions and interconnections that may be referred (e.g., a conversation from yesterday) for details and continuity. This contextualization may further enable avoiding repeating responses, making it more efficient to reference a previous conversation. Yet further, a contextualization module may provide context to the GAIE of other conversations between the user and the system, between other users and the system, and the like.

[2931] In such a contextually maintained instance, a context-enabled GAIE may provide a response regarding forecasted weather that references an earlier period of time. In an example, a context-enabled GAIE may provide a weather-related response such as, "On Monday, we discussed the weather, you asked if you would need an umbrella on Wednesday, and I answered 'probably

not' based on the forecast at that time. I need to inform you that the updated weather forecast indicates that rain may be more likely on Wednesday, so you probably may need an umbrella."

[2932] Other capabilities of emerging GAIE systems may include adapting a GAIE to the generation and operation of digital avatars. In example embodiments, digital avatars may be programmed with their own visual representations. To accomplish greater similarity between an avatar and its owner based on visual and audio interpretation of users, a GAIE training and/or pre-training data set may require information about body language and nonverbal cues, such as gaze, posture, speech pitch and volume, and the like.

[2933] Emerging GAIE systems may include determining and adapting responses with variations and nuances based on, for example, user activities. A user's physical disposition may influence content production by a GAIE (e.g., presenting different cues) based on if the user is sitting, walking, driving, exercising, and the like. Further, a GAIE system may adapt responses to prompts based on variations and nuances of real-life interactions versus voice interfaces versus virtual reality. Other aspects that may impact GAIE responding to prompts may include variations and nuances of different cultures, demographics, and the like. Yet further, in example embodiments, methods and systems for advanced GAIE training and operation may include recognition of higher-level communication features of users (humor, sarcasm, dishonesty, double entendre, etc.) and user emotional state, for example.

[2934] In example embodiments, methods and systems for enhancing GAIE platforms, such as those described herein, may include configuring a GAIE to participate in multi-user dialogue, where strict turn-taking interaction with one person might be difficult in a group setting, where the context of who may be speaking to whom matters for each expression. The more fluid multi-user conversational structure vs. turn-taking structure may indicate advances to a GAIE may include developing understanding of: social interactions and cues, such as to whom each expression may be directed; group dynamics (e.g., who may be the group leader?) and interpersonal relationships; the notion of threaded discussions with branches; concurrent discussions between various sub-groups of a group; when to chime in with input so as to avoid interrupting other users; some notion about conversational balance, to avoid dominating the conversation; tact: users' sensitivity about personal information, and when it may and cannot be shared in a group setting based on context, relationships with other users, and the like.

[2935] Independent of whether interactions are one-on-one or multi-user, it is envisioned that a GAIE may be adapted to evolve beyond a turn-taking paradigm. In an example, a GAIE may currently create media (images, music, video, and the like) based on a user prompt (that itself may be one or more types of media), and may refine the created media based on user interactions, such as changing the content in certain ways or extending the boundaries of an image with more content that may be consistent with the existing content (e.g., outpainting). A more sophisticated version of generative AI may flexibly and continuously adapt its generated content to contextual user input and interactions. In an example, generating media may be adapted by the GAIE in response user integration with the generated media content, such as in response to allowing a user to virtually walk around inside the content to interact with and/or react to content items. Such a media-adapting

GAIE may generate new content or update the content based on the user input/content virtual interactions. Yet further to facilitate a user to virtually interact immersively with generated content details about the user may be considered part of the criteria for newly generating and/or updating the media.

[2936] In example embodiments, a media-output enabled GAIE without user immersive interaction and feedback may generate media (e.g., a first image) based on a prompt in which a user specifies a theme for a story. The user may then specify a series of scenes that follow, and the GAIE generates an image for each scene, leading to a storyboard series for the story.

[2937] When a media-out enabled GAIE is teamed with user immersive capabilities, the user may control, for example, an avatar that may walk around within the scene and interact with generated media objects. Based, for example, on an order and manner with which the user traverses the scene and interacts with the objects, the generative algorithm may generate new content (e.g., the user looks at a particular painting on the wall of a gallery and then opens the curtains of a window). Outside the window may be an entire world that may be consistent with the particular painting that the user viewed. If the user chooses to move the avatar into that world, the painting on the wall updates to reflect the user's interactions.

[2938] In another example of immersive user-generated media content engagement, a user may request a science fiction story. In addition to generating a story based on tropes that are generally relevant to science fiction, the GAIE may include tropes that are likely familiar to the user, such as based on the user's age, culture, other interests, etc. (such as science fiction versions of characters that are well-known in the oeuvre of myth and literature to which the user belongs). In some cases, the algorithm may even include individuals in the created story that are analogous to celebrities or public figures in the user's culture or generation, or even the user's own friends and acquaintances.

[2939] In example embodiments, a superintelligence system may be based on a pre-trained GAIE that facilitates automated discovery of relevant domain-specific knowledge and examples. The superintelligence system may further leverage pre-trained advanced GAIE to leverage domain-specific examples to generate content. Yet further the superintelligence system may include a genetic programming capability to create novel variation. In example embodiments, a superintelligence system may further include feedback systems (e.g., collaborative filtering and automated outcome tracking) to prune variation to favorable outcomes (financial, personalization, group targeting, and the like).

[2940] In example embodiments, a GAIE may be pre-trained for use by and/or in cooperative operation with a digital twin engine, such as an instance of an executive digital twin and the like. In an exemplary deployment, a GAIE may interact with a digital twin to provide a narrative about a topic of the digital twin to give to a viewer. In this example, the digital twin may interact with the GAIE (e.g., through an API and the like) to generate a narrative summary for a CEO and a detailed narrative for a CFO.

[2941] Executive digital twins may be configured for a particular role or user. Therefore, a GAIE system with a digital twin interface may improve executive digital twin capabilities by curating the data for and populating content for consumption by executive digital twins for different roles. In

an example a GAIE may receive information about the executive digital twin as well as about the intended human being represented by the executive digital twin (e.g., the role of the user). The GAIE may determine a degree of narrative detail for each executive digital twin. This may be based on generic executive digital twin/user role criteria and/or refined through interaction with a particular user for the executive digital twin. In example embodiments, a CEO with a tech focus may receive more "in-depth" narrative relating to tech or R&D, whereas a CEO with a financial background may end up receiving narratives that are more focused on financial analysis but less granular on tech-related features.

[2942] In example embodiments, a GAIE system that interacts with a digital twin engine (e.g., an executive digital twin instance and/or engine) may determine of the potential universe of content on which it is trained, what may be relevant and what may be noise or unrelated for the specific narrative topic, the target human consumer, and the like. Based on this relevance determination, the GAIE system may generate the output data based on the relevant data and the determined degree of detail.

[2943] Further, the GAIE system may also select real time data sources to connect to a target / requesting executive digital twin. The GAIE may further configure consumption pipelines for those sources on the spot (e.g., data source identification, data requests for identified data sources, API configuration, and the like). Therefore, in this example the GAIE system would be identifying data sources and connecting them to an executive digital twin instance/engine.

[2944] An example use case may include an executive digital twin that has access to full financial data from a previous time-frame (e.g., a previous year/quarter/month, and the like). The executive digital twin may enable access by the GAIE to all of this data. The GAIE may determine a degree of detail of the data for the intended viewer (e.g., target consumer of a narrative regarding a topic captured in the full financial data).

[2945] In the case of a target consumer/view having a role of CEO, the GAIE may determine that the narrative for the CEO will include key insights but not full details. The GAIE may then generate a narrative of the top insights for a target time-frame (e.g., a current quarter) from at least the received data.

[2946] A pre-trained GAIE may be used to generate, manage, and/or manipulate digital twins, such as by describing attributes of a digital twin, describing interactions with other digital twins or environments, describing simulations, using digital twin simulation data to generate content, enabling context-adaptive executive digital twins, facilitating development of narratives about ongoing, real time operations, tuned to the preferred conversation style of a user represented by a digital twin, and the like. In example embodiments, a context-adaptive executive digital twin integrated with a generative conversational AI system may be configured to generate a set of narratives about operations of an enterprise based on an input data set of real-time sensor data from the operations of the enterprise. The digital twin (or human user) may prompt the GAIE and/or conversational AI system to compare financials with real-time sensor data.

[2947] In example embodiments, a GAIE may be adapted for conversational querying of machine and/or system states, such as in association with an executive digital twin. A generative AI

conversational agent may be an instance of a GAIE that is trained on a structure, function, and operational parameters of a system optionally in an industrial environment. The agent may provide on-demand operational data and insights, including selection of noteworthy operational parameters.

[2948] In example embodiments, a GAIE may be adapted (e.g., pre-trained) for producing qualitative feedback and/or description of a fault prediction based on empirical operational data. A generative AI conversational agent may be configured as an instance of a GAIE that provides predictive notifications of potential fault conditions for a machine or system, such as in an industrial environment.

[2949] Further, such a system may conduct an interactive conversation with a user as to the development of a problem with a machine (e.g., gathering information on how the machine was being used, how the user believes that it performed in the past, what the user perceives to be different, and what the user believes to be the problem). The system may further integrate an explanation of a diagnosis or prognosis of a machine analysis with the user's explanation (e.g.: "the vibration you reported hearing last week was caused by this problem"). This may help to cultivate the user's own intuition about the operation of a machine in addition to addressing an immediate problem. In an example, a generative AI conversation engine may receive an observation of a user associated with a machine through an interactive conversation with the user. The engine may reference, include, and/or interface with a machine diagnostic component that diagnoses a condition of a machine based on the observation of the user.

[2950] A GAIE may be adapted (e.g., pre-trained) to facilitate enhancement of AI training data associated with a digital twin application. In example embodiments, a method may include using an AI conversational agent to create synthetic training data.

[2951] Further in association with digital twin technology, a GAIE may be adapted for summarizing highly granular data for consumption by an executive digital twin. In this regard, an executive digital twin system may include an intelligent agent that receives a set of customization features from a user (e.g., an executive represented by the digital twin) that includes a role of the user within an organization. The intelligent agent may also determine a respective granularity level of a report based on the customization features. In example embodiments, the set of customization features include granularity designations for different types of reports. Yet further, the intelligent agent determines the granularity level of a report based on the role of the user within an organization. Further, the subject matter of the report may be generated based on the role of the user within the organization.

[2952] In example embodiments, a speech-based user interface for customizing a level of specificity for generating executive digital twin reports may be operatively coupled to a customized GAIE that processes the speech into a set of report instructions (and optionally report content) based on aspects of the user(s). An example of a speech-based request that may be processed as described may include, "I'd like an executive-summary level report on predictive maintenance" or "I'd like a detailed report on competitor analysis." The speech-based user interface may respond to such a request by directing a corresponding executive digital twin system to feed a specificity level

for parameters to a generative AI engine (e.g., GAIE) as additional input along with the data. In this example, IoT data from manufacturing facilities may be used in predictive maintenance. A response to a prompt regarding preventive maintenance may be customized with a level of specificity based on target report consumer role(s), such as for an operations-based role. A level of specificity may include what are the costs, when is the maintenance needed by, what may be the predicted downtime, how to offset and/or time the maintenance activity, and the like. For a financial-based role, specificity levels may be adapted to address what may be the disruption going to do for the bottom line in the short term; how does this impact our supply; what may the disruption do to our market-share; will it impact our stock price, and the like.

[2953] When a digital twin may be used to model an individual, a fine-tuned GAIE may be used to coordinate the digital twin with the human for improved fidelity (e.g., when the human behaves or reacts differently than the digital twin predicts, a GAIE may initiate a dialogue with the user to determine why, and the results may be used to update the digital twin model for the individual). Instead of having a human expert occasionally participate in automated digital twin model training (e.g., to correct errors or provide new examples, and the like), a corresponding GAIE may be occasionally querying the user to solicit more information to update the digital twin model of the individual. As an example, a system may include a digital twin that models an individual, and may further include a conversation engine that facilitates determining an update of the digital twin based on a conversation with the individual that is associated with a difference between an action of the individual and a corresponding action prediction by the digital twin.

[2954] In example embodiments, a GAIE system may be configured for use in an automated manufacturing environment. In one example, a user may prepare a descriptive prompt of a desired product to have it 3D printed. The GAIE system may generate a 3D printing set of instructions, such as a configuration of an automated 3D printing machine and a rendering indicative of a result of the 3D printing machine following the instructions. In another example, a user may include a photo/video of product as a prompt along with a request for instructions to 3D print an improved version, such as "I want this bike but I want different tires and I want it to be red."

[2955] Another exemplary use of a pre-trained GAIE may include using user behavioral data to generate guiding recommendations for energy conservation, usage shifting, and the like. In particular, a recommendation system for energy conservation, usage shifting, or optimization may include an integrated generative, conversational AI system that adapts generated output based on user behavior from a user behavior data set.

[2956] In example embodiments, an adapted GAIE may facilitate management of energy resources. An energy resource management system may be enhanced to provide advanced intelligence (e.g., superintelligence) to plan, manage, and/or govern distributed energy resources (DERs) and energy generation, storage, consumption, and transmission facilities. Elements of a superintelligent energy management system may include automated discovery of relevant domain-specific knowledge and examples, generative AI to leverage domain-specific examples to generate content, genetic programming to create novel variation, feedback systems (e.g., collaborative filtering and automated outcome tracking) to prune variation to favorable outcomes (financial,

personalization, group targeting, etc., etc.), and the like. In an example, a superintelligent AI-enabled management system may be configured to manage a plurality of systems of an energy edge platform via automated discovery, generative AI, genetic programming, and feedback systems.

[2957] In example embodiments, a GAIE may be adapted (e.g., trained, pre-trained, and the like) for the field of patents to generate patent claims responsive to being provided a patent disclosure. An enabled GAIE may receive patent claims as a prompt and may generate a supportive patent disclosure therefrom. In example embodiments, an enabled GAIE may be trained to understand a patent structure and a claim structure for a plurality of jurisdictions.

[2958] In example embodiments, a GAIE may be pretrained (e.g., finetuned) with a private instance of an enterprise's intellectual property data (e.g., products, business goals, competitive considerations, core inventive ideas, and the like). In example embodiments, a private instance of enterprise data for patent generation may be configured (e.g., as prompt-response pairs) for finetuning the GAIE instance.

[2959] Beyond patent disclosure and figure preparation, a GAIE may be fine-tuned to generate figures, disclosure from figures, claims from figures, office action responses, evidence of use (EOU) for patent monetizing, preparing a matrix of patent claims across a portfolio, high level landscape search strings, enhancement of search strings, and the like. Finetuning may include preparation of prompt-response sets for a range of IP-related actions, such as patent claim assertion, infringement analysis and discovery, claim (term) acceptance and/or rejection, estimate of claim scope broadness, claim quality, and the like. In example embodiments, an IP-tuned GAIE may be pre-trained with information from proceedings related to infringement cases to understand the likelihood of infringement, and the like.

[2960] GAIE training and IP-integration may facilitate elaboration of broadly stated inventive concepts into disclosure that reflects robust enablement and/or support. In an example, an outline may be an input prompt for the purposes of drafting a patent application (e.g., disclosure, figures, summary, abstract, and optionally claims). A generated result may become a portion of a subsequent prompt along with a description of the general theme, category, focus area and/or other categorization or classification of innovation. In an example, one may describe a transaction environment processing platform and ask for examples of a technical implementation, system, and/or method design, such as: "In the context of a transaction environment processing platform as previously described, what types of hardware and software might be used to implement a governance engine for the transaction environment?"

[2961] Regarding an intellectual property (e.g., patent) monetization-focused development process, a GAIE may facilitate predicting, from a market development view, which domains to select and which categories within domains to emphasize based on the ability to determine where business may be shifting over a longer time (e.g., beyond short-term trends). This may include analyzing historical data and current data for one or more IP domains, optionally in near-real time. An IP-monetization-focused GAIE may tie historical and/or current data to investments and actions having occurred in the IP world for, among other things, patent sales and licensing. An IP-monetizing trained GAIE may also develop particular leads and domain categories with the highest

probability of success based on previous sales and/or licensing and/or where the market may be heading. There may be risk in making these decisions but using a trained GAIE may lower this risk so that these decisions become more predictable in the future, especially with company data increasing and likely accessible through various channels.

[2962] A GAIE may be configured, trained, and/or fine-tuned for a range of functions, including, for example, ingestion of proprietary data, determination of a route, determination of an outcome, approval of release/access to data, making a prediction, pattern recognition, and the like. Yet another example application of a fine-tuned GAIE may include layering of voice and visual commands that may be graduated in sound, volume, or spacing similar to flight avionics, thereby generating scripts for voice over of data and/or presentation material. This may enable the development of synthetic speech technology that generates lifelike (AI-generated) voices for podcasts, slideshows, and professional presentations. This may mitigate needs for hiring a voice artist or using any complex recording equipment (e.g., background noise separation, dubbing, and the like).

[2963] In example embodiments, GAIE systems may be configured for facilitating news delivery from NPC-type avatars to adapt current "clickbait" content to conversationally conveyed world news/happenings. In this example, a metaverse environment may include a news-based GAIE conversation agent configured to conversationally inform users of recent events.

[2964] Further in context of metaverse technology, a generative AI conversational agent may be configured to populate the metaverse.

[2965] Yet further within a context of metaverse technology, a GAIE system may be enabled to augment training data for a customized conversational agent with real-time sensor data sets through collecting information from real-world sensors. In an example, a training data augmentation system may be configured for augmenting training of a conversational agent with data from a real-time sensor data set. Further, a metaverse-associated GAIE system may facilitate augmenting training data for a customized conversational agent with process outcome data. A training data augmentation system may be configured for augmenting training of a conversational agent with process outcome data from a process outcome data set, user behavior data, and the like. In example embodiments, a training data augmentation system based on a GAIE may be enabled (e.g., pre-trained) for augmenting training of a conversational agent with user behavior data from a user behavior data set.

[2966] In example embodiments, application of fine-tuned GAIE systems in the field of governance may facilitate advances in automation of governance, such as governing use of copyrighted material. GAIE-based governance systems may further enhance governing AI training, such as conversational AI training data sets for bias and error, governing conversational AI for contextual appropriateness and other stylistic requirements, and the like. A fine-tuned GAIE system may further improve governing secrecy, such as a progression of what elements of secret, proprietary or confidential information are allowed based on a depth of conversation. Governance may further apply to individuals. Therefore, a governance fine-tuned GAIE system may enhance and/or automate determining a measure of trustworthiness of a user that may be interacting with a

generative conversational AI system. Further a governance fine-tuned GAIE system may enrich governance for a generative AI system, such as determining a measure of trustworthiness of a generative conversational AI system. In general, governance use cases may be expanded further in light of GAIE topic-targeting training capabilities.

[2967] A fine-tuned GAIE system may play a role in systematic risk identification, management. and opportunity mining. GAIE-based risk identification systems may respond to risk-related prompts, such as "What may else might we know and should be paying attention to?" by curating data sets and automating the processes of identification of systemic risks, identifying a set of likely scenarios and the risks and opportunities arising from those scenarios, identifying paths for resolution and recommending resolutions.

[2968] Yet another area of risk identification and/or management may involve security concerns with GAIE systems that are configured to generate computer executable code. At the least relying on computers to write computer code raises questions about what security measures are effective and what measures are able to be circumvented by the AI.

[2969] A further area of risk identification, management and/or opportunity harvesting may apply to copyright material. Automated computer code generation may inadvertently introduce copyrighted material, such as algorithms. A risk-finetuned copyright GAIE may assist in detecting candidate copyright violations in any programmatic code, including machine generated code.

[2970] Risk identification of visual training sets (e.g., images, graphs, and the like) may be enhanced by a fine-tuned GAIE that can process these visual training data sets for authenticity indicators that are coded as non-visual data. This may be similar to tail voltage devices providing messages on the end of sine waves. Visual training sets may be coded with non-visual indicators of authenticity that may be detectable by a fine-tuned GAIE.

[2971] Yet another risk-identification related area includes fraud detection. Integrating customer fraud reporting and questioning into pretraining data may enrich holistic scoring, which may comprise a composite score that bridges customer evidence, transactions, and environmental trends. In an example, an AI based fraud detection system may integrate customer fraud reports and questioning into a training/query data set to produce a holistic scoring system, utilizing a composite score that combines customer evidence, transaction data, and environmental trends to provide a comprehensive approach to fraud detection.

[2972] Imaging applications may benefit from fine-tuned GAIE systems. In example embodiments, optical content (e.g., screen shots and the like) may be processed by machine vision systems so that the GAIE may describe a scene in the optical content using a generative conversational AI agent. In example embodiments, a GAIE may be configured as a first AI/NN sub-system in a Dual Process Artificial Neural Network (DPANN) architecture. Such a DPANN architecture may include, as a second NN sub-system, a formal logic-based and/or fuzzy-based system. Together these DPANN systems may implement learning processes, model management, and the like. In example embodiments, a DPANN architecture may include features that describe building and managing large scale models.

[2973] Referring to Fig. 268, a platform for the application of generative AI 26800 may include a robust task-agnostic next-token prediction artificial intelligence model 26802 that operates to predict a next token given a set of inputs encoded as embedded tokens. A robust task-agnostic next-token prediction AI engine 26802 may include deep learning models, which use multi-layered neural networks to process, analyze, and make predictions with complex data, such as language. An objective of the robust next-token prediction AI engine 26802 may include data science modeling through, among other things, use of topic-specific embeddings, attention mechanisms, and decoder-only transformer models. Capabilities of such an engine 26802 may include a pre-training capability to facilitate configuring next-token prediction for specific subject matter (e.g., marketplace item valuation), a tokenizing capability to facilitate converting complex terms into actionable tokens (e.g., converting compound chemical names into fundamental elements), access to distributed training (e.g., data-parallel training and/or model-parallel training, and the like), few-shot learning to reduce training demand for updates, such as new business intelligence data, and the like. In general the next-token prediction AI engine 26802 may combine large language modeling techniques and decoder-only transformer models to generate powerful foundation models for next-token prediction AI content generation.

[2974] In example embodiments, the next-token prediction AI engine 26802 may be structured with a machine learning (sparse Multi-Layer Perceptron) architecture configured to sparsely activate conditional computation using, for example mixture-of-experts (MoE) techniques. A machine learning architecture may be configured with expert modules that may be used to process inputs and a gating function that may facilitate assigning expert modules to process portion(s) of input tokens. A machine learning architecture may further include a combination of deterministic routing of input tokens to expert modules and learned routing that uses a portion of input tokens to predict the expert modules for a set of input tokens.

[2975] A GAIE 26800 may be trained to operate within a domain, such as written language, computer programming language, subject matter-specific domains (e.g., a software orchestrated marketplace domain), and the like to generate content (constructs) that comply with rules of the domain. In general, a GAIE may generate content for any topic for which the GAIE is trained. So, for example, a GAIE may be trained on a topic of pig farmers and may therefore generate language-based descriptions, images, contracts, breeding guidance, textual output, and the like for any of a potentially wide range of pig farmer sub-topics.

[2976] Adapting a generative AI engine for subject matter-specific applications may include pretraining a next-token prediction AI model-based system through the use of, for example, in-context (e.g., application, domain, topic-specific) examples that are responsive to a corresponding prompt. While the next-token predictive capabilities of the underlying next-token prediction AI engine may remain unaffected by this pre-training, subject matter-specific pre-trained instances may be developed/deployed.

[2977] In example embodiments, a platform for the application of generative AI 26800 may include a set of subject matter-specific pretrained examples and prompts 26804. This set 26804 may be configured by analyzing (e.g., by a human expert and/or computer-based expert and/or

digital twin) information that characterizes various aspects of the domain to generate example prompts and preferred and/or correct responses. Pretraining may also include training the next-token prediction AI engine 26802 by sampling some text (e.g., prompt/response sets) from the set of subject matter-specific pretrained examples and prompts 26804 and training it to predict a next word, object, and/or term. Pretraining may also include sampling some images, contracts, architectures, and the like to predict a next token. These prompt-response sub-sets may facilitate pre-training the prediction AI engine 26802 for predicting a next token (e.g., word, object, image element, and the like) for various aspects.

[2978] When an instance is implemented for textual generation, such a GAIE instance may be referred to as a natural language generation system that constructs words (e.g., from sub-word tokens), sentences, and paragraphs for a target subject and/or domain.

[2979] In example embodiments, real-world instances of the platform 26800 may require ongoing updates to facilitate the platform 26800 being responsive as aspects of a domain (e.g., a business entity in the domain) change, such as business goals change, new products are released, competitors merge, new markets emerge, and the like. In this regard, training the platform 26800 with in-context prompts and examples may be automated and repeated as new data is released for an enterprise to prevent snapshot-in-time data aging-based errors. The platform the application of generative AI 26800 may include an ongoing pre-training module 26828 that processes new and updated content into prompt and/or response sets and interactively iterates through rounds of pre-training. New and updated data and/or information may regularly be found in various subject matter specific information sets, such as: a dataset of medical records (e.g., to assist with medical diagnoses), a dataset of legal documents and court decisions (e.g., to provide legal advice), a release of a new product (e.g., images of the product), or a financial dataset such as SEC filings or analyst reports. In example embodiments, uses of the platform 26800 may include applying the pre-training and optimizing techniques to a range of different domains (e.g., medical diagnosis, business operation, marketplace operation, and the like) to produce a fine-tuned domain specific token-predictive engine including ongoing refinement through (daily) in-context pretraining.

[2980] In example embodiments, an ongoing pre-training module 26828 may work with the next-token prediction AI engine 26802 to update a set of subject matter specific tokens that may be maintained in a subject matter specific instance token storage facility 26808. This subject matter specific instance token storage facility 26808 may be referenced by a subject matter specific instance of the next-token prediction AI engine 26802 during an operational mode (e.g., when processing inputs / prompts). In example embodiments, the platform 26800 may include a plurality of sets of subject matter specific tokens that may be maintained by corresponding ongoing pre-training modules 26828.

[2981] Training, however, may not ensure that the responses to prompts are correct every time. In general, a business entity is likely to be less interested in a tool that provides answers that are probably right and may differ from time to time. A product that can provide accurate responses (e.g., including taking actions) based on what the end-user wants vastly increases the potential use cases and product value. A high level of accuracy and integration with operational systems may

enable such a tool to go beyond just generating new content to be more productive; through integration with workflows, it may facilitate automating workflow actions. In this regard, the platform for the application of generative AI 26800 may also include a pre-training optimizing module 26806 that may work cooperatively with the ongoing pre-training module 26828 to further refine accuracy of responses to prompts for a domain. The pre-training optimizing circuit 26806 may facilitate improved accuracy of in-context responses, task-specific fine-tuning, and for sparse model variants of the platform 26800, enrich few-shot learning capabilities. In example embodiments, fine tuning may further benefit the platform by reducing bias that may be present in the training data. This may be essential to ensure subject matter specific jargon is adapted as training data changes (e.g., in the digital marketing/promotional space, ensure that "influencer" is replaced with "creator"). Further, a pre-training optimizing engine 26806 may provide a wider range of prompts and responses based on user preferences (e.g., speaking styles) to enrich the platform's ability to provide user-centric responses. In example embodiments, user-centric responses may include fine tuning the platform 26800 for different roles in an organization. As an example, when a user in a financial planning role inquires about a business development topic, responses may be directed toward the financial planning role (e.g., as compared to a customer/client inquiry about that topic).

[2982] A platform for the application of generative AI 26800 may be used to produce text-based content for a multi-national entity with employees who speak different languages. While the platform 26800 may be trained (and pre-trained) to operate interactively in a plurality of languages, generating automated content may benefit from use of a neural machine translation module 26810. In example embodiments, a portion of the entity in a first jurisdiction may produce content in a first language and resulting recurring generated output (e.g., types of reports and the like) may be generated in the first language. However, employees who speak a second language may benefit from the type of report when translated into the employee's native language. Therefore, associating the neural machine translation module 26810 with the platform may prove valuable while reducing compute demand for the platform 26800.

[2983] Emerging next-token prediction AI systems feature increasingly adaptable next token prediction capabilities. These capabilities may be further adapted to assist in closed problem set solution prediction, such as allocation of resources, deployment of a robotic fleet and the like. To achieve greater prediction capabilities, a subject matter specific next-token prediction AI-based engine, such as the platform for the application of generative AI 26800, may include a solution-predictive engine 26812 that leverages next-token (e.g., next word) predictive capabilities to predict a most-likely solution to a closed solution-set problem. This may be accomplished optionally through use of sets of problem domain-specific pre-training prompts and examples. Such examples may be adapted for different user preferences. In example embodiments, each user in a closed problem set environment may generate prompts and responses that may enable the platform 26800 to respond to the user based on the user's inquiry style. Alternatively, the solution prediction engine 26812 may adapt a user's prompt and/or configure a prompt based on user

preferences to attempt to deliver responses that are consistent with a user's preferences (e.g., engineering-based responses for an engineer role-user and legal-based responses for a lawyer).

[2984] For more complex analysis and decision making/predicting, a formal logic-based AI system 26814 may be incorporated into and/or be referenced by the subject matter specific platform 26800.

[2985] Further, the basic concepts of next-token prediction of a generative AI engine, such as the platform for subject matter based application of generative AI 26800 may be applied to analyzed expressions of images, audio (e.g., encoded text), video (e.g., sequences of related images), programmatic code (domain-specific text with readily understood rules), and the like. Therefore, a next-token prediction AI platform (e.g., platform 26800) may further include an image/video analysis engine 26816 (optionally NN-based) that adds a spatial aspect to the next-token predictive capabilities of a next-token prediction AI system. Images used for training may include 3D CAD images (for a domain that includes physical devices such as vehicles), radiologic images (for a medical analysis domain), business performance graphs, schematics, and the like. In example embodiments, aspects of the underlying task-agnostic next-token prediction AI engine 26802 may be adapted (e.g., different embeddings, neural network structures and the like) for different input formats, such as images, temporal-spatial content, and the like.

[2986] The platform 26800 may further include an expert review and approval portal 26818 through which an expert (e.g., human / digital twin, and the like) can review, edit, and approve content generated. Examples include review and adaptation by a subject matter specific data story expert; a data scientist, and the like. The expert review and approval portal 26818 may operate cooperatively with, for example, the pre-training optimizing module 26806 that may receive and analyze expert feedback (e.g., edits to the content and the like) for opportunities to further optimize the platform 26800.

[2987] The platform 26800 may further include a training data generation facility 26820 that may generate natural language prompts, such as subject matter specific prompts that may be applied by, for example, the pre-training optimizing engine 26806 to increase platform response accuracy and/or efficiency while fine tuning a subject matter specific instance.

[2988] In example embodiments, the platform 26800 may further be configured to access a corpus of domain and/or problem relevant content as a step in responding to a prompt. In example embodiments, the platform may be pre-trained on the content of the corpus. While the content of the corpus may not be directly included in the response, such as if it provides a level of detail beyond what the platform 26800 has been trained to provide in a response, it may be cited in the response to facilitate identifying and expressing sources from which a response is derived. These external source references may be handled via a citation module 26822.

[2989] Business decisions are often context-based. Understanding both the context for a decision and aspects and/or assumptions of the decision process may prove highly valuable for evaluating, for example, competing decisions and/or recommendations. Context may include both tangible and intangible factors. An intangible factor may include historical interactions between parties involved in the evaluation process, for example. A decision process may include not only assumptions on which a decision or recommendation is based, but also criteria by which tangible

factors are processed, evaluated, analyzed, and the like. To provide such context for generated output of the platform 26800, an interpretability engine 26824 may be incorporated into and/or be accessible to the platform 26800. An objective of use of the interpretability engine 26824 may be to generate additional content that reflects context for, among other things, how the next-token prediction AI instance operates and/or generates a corresponding output.

[2990] In example embodiments, the next-token predictive capabilities of a next-token prediction AI engine 26802 may be utilized for developing a set of emergent data science predictive and/or interpretive skills. While such a platform may be trained directly on various data sets, context for elements and results in such data sets may be a rich source of complementary training data. By associating data elements with descriptions thereof, the platform 26800 may gain data science capabilities, such as to group by or pivot categorical sums, infer feature importance, derive correlations, predict unseen test cases, and the like. In this regard, a data science emergent skill development system 26826 may be utilized by the platform to enhance further subject matter specific applicability and utility.

Graph Neural Networks and Processing of Graph Data Using Machine Learning Algorithms

[2991] In various embodiments, one or more techniques involve the processing of graph data using one or more machine learning algorithms. In some such embodiments, the one or more machine learning algorithms include one or more graph neural networks (GNNs). The following discussion provides an overview of graph data and graph neural networks.

[2992] In a graph data set, a set of nodes is interconnected by one or more edges that respectively represents a relationship among two or more connected nodes. In many graph data sets, each edge connects two nodes. In other graph data sets that represent hypergraphs, a hyperedge can connect three or more nodes. In various graph data sets, each of the one or more edges is directed or undirected. An undirected edge represents a relationship that relates two or more nodes without any particular ordering of the related nodes. A first undirected relationship that connects a first node N1 and a second node N2 may be equivalent to a second undirected relationship that also connects a first node N1 and a second node N2. In some such graphs, the relationship represents a group to which the two or more related nodes belong. In some such graphs, the relationship represents an undirected and/or omnidirectional connection between two or more nodes. For example, in a graph representing a geographic region, each node may represent a city, and each edge may represent a road that connects two or more cities and that can be traveled in either direction. By contrast, a directed edge includes a direction of the relationship between a first node and a second node. For example, in a graph representing a genealogy or lineage, each node represents a person, and each edge connects a parent to a child. A first directed edge that connects a first node N1 to a second node N2 is not equivalent to a second directed edge that connects the second node N2 to the first node N1. Some graph data sets include one or more unidirectional edges, that is, an edge with one direction among two or more connected nodes. Some graph data sets include one or more multidirectional edges, that is, an edge with two or more directions among the two or more connected nodes. Some graph data sets may include one or more undirected edges,

one or more unidirectional edges, and/or one or more multidirectional edges. For example, in a graph representing a geographic region, each node may represent a city; one or more unidirectional edges may represent a one-way road that connects a first city to a second city and can only be traveled from the first city to the second city; and one or more bidirectional or undirected edges that represent a bidirectional road between the first city and the second city that can be traveled in either direction. Some graph data may include, for two or more nodes, a plurality of edges that interconnect the two or more nodes. For example, a graph data set representing a collection of devices may include nodes that respectively correspond to each device of the collection and edges that respectively correspond to an instance of communication and/or interaction among two or more of the devices. In such a graph data set, a particular subset of two or more devices may engage in a plurality, including a multitude, of instances of communication and/or interaction, and may therefore be connected by a plurality, including a multitude, of edges.

[2993] Some directed and/or undirected graph data sets may include one or more cycles. For example, in a graph representing a social network, a first edge E1 may connect a first node N1 (representing a first person) and a second node N2 (representing a second person) to represent a relationship between the first person and the second person. A second edge E2 may connect the second node N2 and a third node N3 (representing a third person) to represent a relationship between the second person and the third person. A third edge E3 may connect the third node N3 and the first node N1 to represent a relationship between the third person and the first person. Such cycles can occur in undirected graphs (e.g., edges in a social network graph that indicate mutual relationships among two or more individuals), directed graphs (e.g., edges in a social network graph that indicate that a first person is influenced by a second person, a second person is influenced by a third person, and a third person is influenced by the first person), and/or hypergraphs (e.g., cycles of relationships among three or more clusters that respectively include three or more nodes). Some cyclic graphs may include one or more cycles that are interlinked (e.g., one or more nodes and/or edges that are included in two or more cycles). Other directed and/or undirected graph data sets may be acyclic (e.g., graphs in which nodes are strictly arranged according to a top-down hierarchy). Still other directed and/or undirected graph data sets may be partially acyclic (e.g., mostly acyclic) but may include one or more cycles among one or more subsets of nodes and/or edges.

[2994] In some graph data sets, one or more nodes includes one or more node properties. For example, in a graph representing a geographic area, each node may represent a city, and each node may include one or more node properties that correspond to one or more properties of the city, such as a size, a population, or a latitude and/or longitude coordinate. Each node property may be of various types, including (without limitation) a Boolean value, an integer, a floating-point number, a set of numbers such as a vector, a string, or the like. In some graph data sets, one or more nodes does not include a node property. For example, in a graph data set representing a set of particles, each particle may be identical to each other particle, and there may be no specific data that distinguishes any particle from any other particle. Thus, the nodes of the graph data set may not include any node properties.

[2995] In some graph data sets, one or more edges includes one or more edge properties. For example, in a graph representing a geographic area, each edge may represent a road, and each edge may include one or more edge properties that correspond to one or more properties of the road, such as a distance, a number of lanes, a direction, a speed limit, a volume of traffic, a start latitude and/or longitude coordinate, and/or an ending latitude and/or longitude coordinate. In some graph data sets, a direction of an edge may be represented as an edge property. Alternatively or additionally, in some graph data sets, a direction of an edge may be represented separately from one or more edge properties. In some graph data sets, one or more edges does not include an edge property. For example, in a graph data set representing a line drawing of a set of points, each edge may represent a line connecting two points, and the edges may be significant only due to connecting two points. Thus, the edges of the graph data sets may not include any edge properties.

[2996] In some graph data sets, the graph includes one or more graph properties. Such graph properties may be global graph properties that correspond to one or more properties of the entire graph. For example, in a graph data set representing a geographic region, the graph may include graph properties such as a total number of nodes and/or cities, a two-dimensional or three-dimensional area represented by the graph, and/or a latitude and/or longitude of a center of the graph. Such graph properties may be global graph properties that correspond to one or more properties of all of the nodes of the graph. For example, in a graph data set representing a geographic region, the graph may include graph properties such as an average population size of the cities represented by the nodes and/or an average connectedness of each city to other cities included in the graph.

[2997] Some graph data sets include a single set of data that includes all nodes and all edges. For example, a graph representing a geographic region may include a set of nodes that represent all cities in the geographic region. Some other graph data sets include one or more subgraphs, wherein each subgraph includes a subset of the nodes of the graph and/or a subset of the edges of the graph. For example, a graph representing a geographic region may include a number of subgraphs, each representing a subregion of the geographic region, and the edges that interconnect the cities within each subregion. As another example, a graph representing a geographic region may include a first subgraph representing cities (e.g., groups of people over a threshold population size and/or population density) and a second subgraph representing towns (e.g., groups of people under the threshold population size and/or population density). In some graph data sets, each node and/or each edge belongs exclusively to one subgraph. In some graph data sets, at least one node and/or at least one edge can belong to two or more subgraphs. For example, in a graph representing a geographic region that includes a number of subgraphs respectively representing different geographic subregion, each node representing a city may be exclusively included in one subgraph, while each edge may interconnect two or more cities within one subgraph (i.e., within one subregion) or may interconnect a first city in a first subgraph (i.e., within a first subregion) and a second city in a second subgraph (i.e., within a second subregion).

[2998] Graph neural networks can include features and/or functionality that are the same as or similar to the features and/or functionality of other neural networks. For example, graph neural

networks include one or more neurons arranged in various configurations. Each neuron receives one or more inputs from the graph data set or another neuron, evaluates the one or more inputs (e.g., via an activation function), and generates one or more outputs that are delivered to one or more other neurons and/or as an output of the graph neural network. Examples of activation functions that can be included in various neurons of the graph neural network include (without limitation) a Heaviside or unit step activation function, a linear activation function, a rectified linear unit (ReLU) activation function, a logistic activation function, a tanh activation function, a hyperbolic activation function, or the like.

[2999] As an example, some graph neural networks include only a single neuron, or only a single layer of neurons that is configured to receive graph data as input and to provide graph data as output of the graph neural network. Some graph neural networks are arranged in a series of two or more layers, wherein input is received by neurons included in a first layer. The output of one or more neurons included in the first layer is delivered, as input, to one or more neurons included in a second layer. For example, each neuron in the first layer may include one or more synapses that respectively interconnect the neuron to one or more neurons of the second layer. In many graph neural networks, each neuron N1 of a preceding layer L1 is connected to each neuron N2 of a following layer by a synapse that includes a weight W. Neuron N2 receives, as input, the output of the neuron N1 multiplied by the weight of the synapse connecting neuron N1 and neuron N2. In many neural networks, layer L1 includes a bias B, which is added to the product of the output of neuron N1 and the weight W of the synapse connecting neuron N1 and neuron N2. As a result, the input to neuron N2 includes the sum of the bias B of layer L1 and the product of the output of neuron N1 and the weight W of the synapse connecting neuron N1 and neuron N2. The output of the neurons included in the second layer can be provided as an output of the graph neural network and/or as input to one or more neurons included in a third layer. Each layer of the graph neural network may include the same number of neurons as a preceding and/or following layer of the graph neural network, or a different number of neurons as preceding and/or following layer of the graph neural network.

[3000] As another example, some graph neural networks include one or more layers that perform particular functions on the output of neurons of another layer, such as a pooling layer that performs a pooling operation (e.g., a minimum, a maximum, or an average) of the outputs of one or more neurons, and that generates output that is received by one or more other neurons (e.g., one or more neurons in a following layer of the graph neural network) and/or as an output of the graph neural network. For example, some graph neural networks (e.g., graph convolution networks) include one or more convolutional layers, each of which performs a convolution operation to an output of neurons of a preceding layer of the graph neural network.

[3001] As another example, some graph neural networks include memory based on an internal state, wherein the processing of a first input data set causes the graph neural network to generate and/or alter an internal state, and the internal state resulting from the processing of one or more earlier input data sets affects the processing of second and later input data sets. That is, the internal state retains a memory of some aspects of earlier processing that contribute to later processing of

the graph neural network. Examples of graph neural networks that include memory features and/or stateful features include graph neural networks featuring one or more gated recurrence units (GRUs) and/or one or more long-short-term-memory (LSTM) cells.

[3002] As another example, some graph neural networks feature recurrent and/or reentrant properties. For example, at least a portion of output of the graph neural network during a first processing is included as input to the graph neural network during a second or later processing, and/or at least a portion of an output from a layer is provided as input to the same layer or a preceding layer of the graph neural network. As another example, in some graph neural networks, an output of a neuron is also received as input by the same neuron during the same processing of an input and/or a subsequent processing of an input. The output of the neuron may be evaluated (e.g., weighted, such as decayed) before being provided to the neuron as input. As another example, some graph neural networks may include one or more skip connections, in which at least a portion of an output of a first layer is provided as input to a third layer without being processed by a second layer. That is, the output of the first layer is provided as input both to the second layer (which generates a second layer output) and to the third layer. In some such graph neural networks, the third layer receives, as input, either the output of the first layer or the output of the second layer. That is, the third layer multiplexes between the output of the first layer and the output of the second layer. Alternatively or additionally, in some such graph neural networks, the third layer receives, as input, both the output of the first layer and the output of the second layer (e.g., as a concatenation of the output vectors to generate the input vector for the third layer), and/or an aggregation of the output of the first layer and the output of the second layer (e.g., a sum or average of the output of the first layer and the output of the second layer). Examples of graph neural networks that include one or more skip connections include jump knowledge networks and highway graph neural networks (highway GNNs).

[3003] As another example, some graph neural networks include two or more subnetworks (e.g., two or more graph neural networks that are configured to process graph data concurrently and/or consecutively). Some graph neural networks include, or are included in, an ensemble of two or more neural networks of the same, similar, or different types (e.g., a graph neural network that outputs data that is processed by a non-graph neural network, Gaussian classifier, random forest, or the like). For example, a random graph forest may include a multitude of graph neural networks, each configured to receive at least a portion of an input graph data set and to generate an output based on a different feature set, different architectures, and/or different forms of processing. The outputs of respective graphs of the random graph forest may be combined in various ways (e.g., a selection of an output based on a minimization and/or maximization of an objective function, or a sum and/or averaging of the outputs) to generate an output of the random graph forest.

[3004] In these and other graph neural networks, the number of layers and the configuration of each layer of the graph neural network (e.g., the number of neurons and the activation function used by each neuron of each layer) can be referred to as hyperparameters of the graph neural network that are determined upon generation of the graph neural network. The weights of node synapses and/or the biases of the layers can be referred to as parameters of the graph neural network

that are learned through a training or retraining process. Further explanation and/or examples of various concepts of other types of neural networks that can also apply to graph neural networks, and additional concepts that apply to other types of neural networks that can also be included in graph neural networks, are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

[3005] Unlike other types of neural networks, graph neural networks are configured to receive, process, generate, and/or transform one or more graph data sets. Some graph neural networks are configured to receive data representing and/or derived from a graph data set, such as an input vector that includes data representing one or more nodes of the graph (optionally including one or more node properties of one or more nodes), one or more edges of the graph (optionally including one or more edge properties of one or more edges), and/or one or more graph properties of the graph. Some graph neural networks are configured to receive an input vector comprising all of the data of a graph data set (e.g., all of the data representing all nodes, all edges, and the graph). Some graph neural networks are configured to receive an input vector comprising only a portion of the data of a graph data set (e.g., only a subset of the nodes of the graph and/or only a subset of the edges of the graph). For example, some graph data sets include a number of subgraphs, and the input vector to the graph neural network includes the data for all of the nodes and/or all of the edges included in one subgraph of the graph. The entire graph can be processed by processing (e.g., concurrently and/or consecutively) each subgraph and combining the output resulting from the processing of each subgraph. As another example, a graph data set representing a set of users of a social network may be processed by a graph neural network that receives, as input, a subset of nodes that correspond to the most influential users of the social network (e.g., those having more than a threshold number of social network connections) and a subset of edges that interconnect the nodes representing those users. Some graph neural networks are configured to receive, as input, data derived from a graph neural network. For example, a graph data set representing a social network may be processed by a graph neural network that receives, as input, data associated with messages exchanged among users of the social network, and provides, as output, an analysis of the messages. Some graph neural networks are configured to receive, as input, non-graph data (e.g., an input vector including coordinates of roads and/or cities in a geographic region) and generate graph data as output (e.g., a graph including nodes that represent the cities, and edges that represent roads interconnecting the nodes representing the cities).

[3006] Some graph neural networks are configured to process input data as graph data. As an example, some graph neural networks are configured to receive, as input, data that represents each of one or more nodes of a graph data set and one or more edges that respectively interconnect two or more nodes of the graph data set. The graph neural network may process a state of each node and/or edge of the input graph data in order to generate an updated state of the node and/or edge. The term "message passing" refers to evaluating and updating the state of a node N or an edge E of a graph based on the states of one or more neighboring nodes N and/or connecting edges E. For example, for each node N1, the graph neural network may evaluate the state of node N1 and/or states of a set of nodes N that are connected to node N1 by at least one edge (e.g., a neighborhood

of nodes that includes N1) and may determine an updated state of node N1 based on the state of the node N1 and/or the states of the neighboring nodes N. As another example, for each node N1, the graph neural network may evaluate the state of the node N1 and/or the states of a set of edges E that connect node N1 to one or more other nodes of the graph, and may determine an updated state of node N1 based on the state of the node N1 and/or the states of the edges E. As yet another example, for each edge E1 of the input graph, the graph neural network may evaluate a state of the edge E1 and/or the states of a set of nodes N of the graph that are connected to the edge E1 and may determine an updated state of edge E1 based on the state of the edge E1 and/or the states of the connected nodes N. As yet another example, for each edge E1 of the input graph that connects a set of nodes N of the graph, the graph neural network may evaluate the state of the edge E1 and the states of the set of edges E that are also connected at least one of the set of nodes N and may determine an updated state of edge E1 based on the state of the edge E1 and/or the states of the other edges. In these and other scenarios, each node N and/or each edge E is evaluated and updated based on a collection of "messages" corresponding to the states of neighboring nodes N and/or connecting edges E.

[3007] In some graph neural networks, each node N1 is updated based a neighborhood of size 1, including only on the states of the edges E that are directly connected to node N1 and/or the states of the other nodes N that are directly connected to node N1 by an edge. In some other graph neural networks, each node N1 is updated based a neighborhood of a size S greater than 1, including the states of other nodes N that are within S edge connections of node N1 and/or edges E that are connected to any such nodes N. In some graph neural networks, each edge E1 is updated based a neighborhood of size 1, including only on the states of the nodes N that edge E1 connects and/or the edges E that are also connected to the nodes N that edge E1 connects. In some other graph neural networks, each edge E1 is updated based a neighborhood of a size S greater than 1, including the states of other nodes N that are within S edge connections of node N1 and/or the set of edges E that are connected to any such nodes N. In some graph neural networks with a neighborhood of size greater than 1, one or more first layers of neurons process each node and/or edge based on the nodes and/or edges within a neighborhood of size 1; a second one or more following layers of neurons further process each node and/or edge based on the nodes and/or edges within a neighborhood of size 2; and so on. That is, the first one or more layers update the state of each node and/or edge based on the states of the directly connected nodes and/or edges, and each following one or more layers further updates the state of each node and/or edge additionally based on the states of indirectly connected nodes and/or edges that are one or more further connections.

[3008] In some graph neural networks, the states of nodes N and/or edges E are evaluated and updated concurrently (e.g., the graph neural network may evaluate the features relevant to each node N and/or each edge E to determine an update, and may do so for all nodes N and/or all edges E, before applying the updates to update the internal states of each node N and/or each edge E). In some graph neural networks, the states of nodes N and/or edges E are evaluated and updated consecutively (e.g., the graph neural network may evaluate the features relevant to a first node N1 and update the state of node N1 before evaluating the features relevant to a second node N2 and

updating the state of node N2). In some graph neural networks, the states of the nodes N and/or the edges E are consecutively evaluated and updated according to a sequential order (e.g., the graph neural network first evaluates and updates a state of a first node N1 that is of a high priority, and then evaluates and updates a state of a first node N2 that is of a lower priority than N1). In some graph neural networks, a state of a node N2 is evaluated after updating a state of a node N1 and, further, based on the updated state of node N1. In some graph neural networks, a state of an edge E2 is evaluated after updating a state of an edge E1 and, further, based on the updated state of an edge E1. In some graph neural networks, the states of nodes N are concurrently evaluated and updated, and then the states of edges E are concurrently evaluated and updated concurrently. In some graph neural networks, the states of edges E are concurrently evaluated and updated, and then the states of nodes N are concurrently evaluated and updated concurrently. These variations in the order of updating the nodes N and/or edges E can be variously combined with the previously discussed variations in the processing of neighborhoods. For example, a graph neural network may include a first one or more layers that are configured to evaluate and concurrently update the states of all nodes and edges within a neighborhood of size 1, followed by a second one or more layers that are configured to evaluate and concurrently update the states of all nodes and edges within a neighborhood of size 2. Another graph neural network may a graph neural network may include a first one or more layers that are configured to evaluate and concurrently update the states of all nodes within a neighborhood of size 1, followed by a second one or more layers that are configured to evaluate and concurrently update the states of all nodes within a neighborhood of size 2, further followed by one or more layers that are configured to update the states of all edges within a neighborhood of size 1 or more.

[3009] Some graph neural networks are configured to evaluate and/or update one or more node properties of one or more nodes of a graph data set. For example, a graph representing a social network may include nodes that represent people, and a graph neural network may evaluate the nodes and/or edges of the graph to predict one or more node properties that correspond to attributes of the person, such as a type of the person, an age of the person, or an opinion of the person. Some graph neural networks are configured to evaluate and/or update one or more edge properties of one or more edges of a graph data set. Some graph neural networks are configured to evaluate and/or update one or more edge properties of one or more edges of a graph data set. For example, a graph representing a social network may include nodes that represent people and edges that represent relationships between people, and a graph neural network may evaluate the nodes and/or edges of the graph to predict one or more edge properties that correspond to attributes of a relationship among two or more people, such as a type of the relationship, a strength of a relationship, or a recency of the relationship. Some graph neural networks are configured to evaluate and/or update one or more graph properties of the graph data set. For example, a graph representing a social network may include nodes that represent people and edges that represent relationships between people, and a graph neural network may evaluate the nodes and/or edges of the graph to predict a feature of a social group to which all of the people belong, such as a common interest or a common demographic trait that is shared by at least many of the people of the social network.

[3010] Some graph neural networks are configured to generate graph data as output. The generated graph data may include one or more nodes (optionally including one or more node properties), one or more edges (optionally including one or more edge properties), and/or one or more graph properties. The generated graph data may be based on input graph data. Some graph neural networks may be configured to receive at least a portion of a graph data set as input, and may generate, as output, modified graph data. As an example, the input graph data set may include a number of nodes and a number of edges interconnecting the nodes, and in the output graph data set generated by the graph neural network, each of the nodes and/or edges of the graph may have been updated based on one or more nodes and/or one or more edges of the input graph data. For example, an input graph data set may represent a social network including nodes representing people and edges representing relationships between people. A graph neural network may be configured to receive at least a portion of the input graph data set, and may output an adjusted graph data set, wherein a state at least one of the nodes and/or at least one of the edges is updated based on the processing of the input data set. For example, various edges representing relationships may be updated to include additional data (e.g., edge properties) to represent an updated relationship between two people represented by nodes. Various nodes may be updated to include additional data (e.g., node properties) to represent updated information about corresponding people based on the relationships. Various graph properties of at least a portion of the graph data set may be updated based on the updated edges and/or nodes, e.g., a new common interest that is shared among many of the people in the social network.

[3011] Some graph neural networks may be configured to output graph data that includes one or more newly discovered nodes based on the input graph data set. For example, an input graph data set representing travel events may include edges that include routes of travelers and nodes that represent locations of interest. A graph neural network may receive the input graph data set, and based on processing of the routes of the travelers, may output an updated graph data set that includes a new node that represents a new location of interest (e.g., a destination of a large number of recent travelers). The output of the graph neural network may include, for one or more new or existing nodes, one or more new or updated node properties (e.g., a classification of the location of interest based on the travel routes). Alternatively or additionally, some graph neural networks may be configured to output graph data that excludes one or more existing nodes of an input graph data set. For example, based on processing the input data set representing routes of travelers, a graph neural network may output an updated graph data set that excludes one of the nodes of the input graph data set representing a location that is no longer a location of interest (e.g., a destination that travelers no longer visit).

[3012] Some graph neural networks may be configured to output graph data that includes one or more newly discovered edges based on the input graph data set. For example, an input graph data set may represent a social network including nodes that represent people and edges that represent connections between people. A graph neural network may receive the input graph data set, and based on processing of the people and connections, may output an updated graph data set that includes a new connection between two people (e.g., a likely relationship based on shared traits

and/or mutual relationships with a number of other people representing a social circle). The output of the graph neural network may include, for one or more new or existing edges, one or more new or updated edge properties (e.g., a classification of a relationship between two or more people). Alternatively or additionally, some graph neural networks may be configured to output graph data that excludes one or more existing edges of an input graph data set. For example, based on processing the input data set representing a social network, a graph neural network may output an updated graph data set that excludes one or more of the edges of the input data set representing a relationship that no longer exists (e.g., a lost connection based on a splitting of a social circle).

[3013] Some graph neural networks may output graph data that is based on data that does not represent an input graph data set. For example, a graph neural network may be configured to receive non-graph data, such as lists of travel routes of drivers, and may generate and output a graph data set including nodes that represent locations of interest and edges that interconnect the locations of interest. Conversely, some graph neural networks may receive input that includes at least a portion of a graph data set and that outputs non-graph data based on the input graph data. For example, a graph neural network may be configured to receive input including graph data, such as a graph of a social network including nodes that represent people and edges that represent connections, and to output non-graph data based on analyses of the input graph data, such as statistics about the people represented in the social network and activity occurring therein.

Properties

[3014] Graph neural networks, including (without limitation) those described above, may be subject to various properties and/or considerations of design and/or operation. These considerations may affect their architecture, processing, implementation, deployment, efficiency, and/or performance.

[3015] As previously discussed, graph neural networks may include edges with varying directionality, such as undirected edges (e.g., edges that represent distances between pairs of nodes that represent cities in a graph that represents a region), unidirectional edges (e.g., edges that represent parent/child relationships among nodes that represent people in a graph that represents a genealogy or lineage), and/or multidirectional edges (e.g., bidirectional edges that represent bidirectional roads between nodes that represent cities in a graph that represents a region). In some graph data sets, all of the edges have the same directionality (e.g., all edges are undirected). A graph neural network can be configured to receive an input vector corresponding to the input data set and to process the edges according to the uniform directionality of the edges (e.g., processing undirected edges without regard to the order in which the nodes are represented as being connected to the edge). Other graph data sets may include edges with different directionality (e.g., in a graph that represents a region, edges can represent roads between nodes that represent cities, and each edge can be either unidirectional to represent a one-way road or bidirectional to represent a two-way road). A graph neural network can be configured to receive an input vector corresponding to the input data set and to process the edges according to the distinct directionality of each edge (e.g., processing a unidirectional edge in a different manner than a bidirectional edge). As one such example, the graph neural network can interpret a bidirectional edge connecting two nodes N1, N2

as a first unidirectional edge that connects node N1 to N2 and a second unidirectional edge that connected node N2 to node N1. The pair of unidirectional edges can share various edge properties and/or can be evaluated and/or updated in a same or similar manner (e.g., for a pair of unidirectional edges corresponding to a bidirectional road, the graph neural network can process data representing a weather condition in a same or similar manner to both unidirectional edges associated with the bidirectional road).

[3016] As previously discussed, some graph neural networks are configured to process nodes according to a "message passing" paradigm, in which the evaluation of each node N1 is based on the states and/or evaluations of other nodes within a neighborhood of the node N1 and/or the edges that connect the node N1 to other nodes in the neighborhood of the node N1. That is, the state of each node in the neighborhood of the node N1 and/or the state of each edge that connects N1 to other nodes of the neighborhood serves as a "message" that informs the evaluation and/or updating of the state of node N1 by the graph neural network. Alternatively or additionally, the evaluation of each edge E1 is based on the states and/or evaluations of other edges within a neighborhood of the edge E1. That is, the state of each node connected by edge E1, and, optionally, the states of other nodes connected to those nodes and/or other edges in such connections, serves as a "message" that informs the evaluation and/or updating of the state of edge E1 by the graph neural network. In each case, the size of the neighborhood can vary; for example, the graph neural network can evaluate each node according to a one-hop neighborhood or a multi-hop neighborhood. Graph neural networks that perform multi-hop neighborhood evaluation can include multiple layers, where a first one or more layers are configured to process a first hop between a node N1 and a one-hop neighborhood including its directly connected neighbors and/or directly connected edges, and a second one or more layers following the first one or more layers are configured to process a second hop between the nodes and/or edges of the one-hop neighborhood and additional nodes and/or edges that are directly connected to the nodes and/or edges of the one-hop neighborhood. In this manner, each node N1 is first evaluated and/or updated based on message passing among the one-hop neighborhood, and is then evaluated and/or updated based on additional messages within the two-hop neighborhood, etc. Other architectures of graph neural networks may perform multi-hop neighborhood evaluation in other ways, e.g., by processing individual clusters of nodes and/or edges to perform message passing among the nodes and/or edges of each cluster, and then performing additional message passing between clusters to update nodes and/or edges of each cluster based on the nodes and/or edges of one or more neighboring clusters.

[3017] In some scenarios, a graph may include nodes and/or edges that are stored, represented, and/or provided as input that is not subject to any particular order (e.g., nodes representing points in a line drawing may not have any node properties, and may therefore be represented in arbitrarily different orders in the input graph data set). In such scenarios, a multitude of semantically equivalent input graph data sets may be logically equivalent to one another. That is, a first representation of a graph may include the nodes and/or edges in a particular order, while a second representation of the same graph may include the same nodes and/or edges in a different order. While both representations of the graph are logically equivalent, the different ordering in which

the nodes and/or edges are provided as input to the graph neural network may cause the graph neural network to provide different output. In other scenarios, a graph comprising a set of nodes and a set of interconnecting edges may be organized, stored, and/or represented in a particular order. For example, the nodes may be ordered according to a property of the nodes, and/or edges may be ordered according to a property of the edges (e.g., in a social network, nodes representing people may be ordered according to the alphabetical order of their names, and edges representing relationships may be ordered according to the alphabetical order of the names of the related people). In such scenarios, changes to the order and/or the selected subsets of graph data may result in different input data sets that represent the same or similar (e.g., logically equivalent) graphs. Due to the manner in which a graph neural network processes the input graph data set, logically equivalent input graph data sets may result in different and logically distinct output data.

[3018] In such scenarios, it may be undesirable for the graph neural network to generate different output for different but logically equivalent representations. That is, it may be desirable for the graph neural network to provide the same or equivalent output for different but logically equivalent representations of a graph. Graph neural networks that exhibit this property can be referred to as "permutation invariant," that is, capable of providing output that does not vary across permutations in the representation of the input graph data set. A variety of techniques may be used to achieve, improve, and/or promote permutation invariance. Some such techniques involve changing representations of the input data set. For example, before processing an input graph data set, the graph neural network may reorder the input data set (e.g., by reordering the units of an input vector) such that nodes and edges are represented in a consistent order. As one such example, an input graph data set may include nodes that represent cities, and the input graph data set may include the nodes and/or edges in varying orders. Prior to processing the input graph data set, the graph neural network may reorder the nodes based on latitude and longitude coordinates of the cities, and the edges can similarly be reordered based on the latitude and longitude coordinates of the nodes connected by each edge. Thus, any representation of the graph including nodes that represent the same set of cities is processed in a similar manner. Similar reordering may involve various node properties and/or edge properties, including (without limitation) an alphabetic ordering of names in a graph including nodes that represent people, a chronological ordering of dates in a graph including nodes that represent events, a numeric ordering of content-based hashcodes in a graph including nodes that represent objects, and/or a numeric ordering of identifiers in a graph including nodes that possess unique numeric identifiers. Other techniques for achieving, improving, and/or promoting permutation invariance involve transforming an input graph data set into a different, permutation-invariant representation that is provided as input to and processed by the graph neural network. For example, a graph data set representing a two-dimensional image or a three-dimensional point cloud may include nodes that represent pixels and edges that represent spatial relationships (e.g., distances and/or orientations) between respective pairs of pixels of the image or respective pairs of points in the point cloud. Different orderings of the pixels and/or points may result in differently ordered, but logically equivalent, graph data sets for a particular image or point cloud. Instead of processing the graph data sets as input, a graph neural network may be configured

to convert the input graph data set into a spectral representation, e.g., based on a spectral decomposition of a Laplacian L of the input graph data set. Instead of encoding information about individual pixels and/or points, the spectral representation instead encodes spectral components of the input graph data sets. The spectral components can be ordered in various ways (e.g., by frequency and/or polynomial order) to generate a permutation-invariant input vector, and the processing of the permutation-invariant input vector by a graph neural network may result in invariant (e.g., identical or at least similar) output of the graph neural network for various permutations of the input graph data set.

[3019] Alternatively or additionally, some techniques for achieving, improving, and/or promoting permutation invariance may relate to the structure of the graph neural network. For example, as an alternative or addition to reordering an input graph data set, a graph neural network may include one or more layers of neurons that process an input vector and generate permutation-invariant output. As one such example, a graph neural network may include a pooling layer that receives an input vector (e.g., an input vector corresponding to an input graph data set, and/or an input vector corresponding to an output of one or more previous layers of the graph neural network) and generates output that is pooled over the input, such as a minimum, maximum, or average of the units of the input. Because operations such as a minimum, maximum, and/or average over a data set are permutation-invariant mathematical operations, the graph neural network may therefore exhibit permutation-invariance of output based on the pooling operation for differently ordered but logically equivalent representations of a particular graph data set. As another such example, a graph neural network may include a filtering layer that receives an input vector (e.g., an input vector corresponding to an input graph data set, and/or an input vector corresponding to an output of one or more previous layers of the graph neural network) and generates output that is filtered based on certain permutation-invariant criteria. For example, in a graph representing a social network that includes nodes representing people, a layer of the graph neural network may filter the nodes to limit the input data set based on the top n nodes of the graph neural network that correspond to the most influential people in the social network. Such filtering may be based, e.g., on a count of the edges of each node (i.e., a count of the number of relationships of each person to other people of the social network), or a weighted calculation based on the influence of the nodes to which each node is related and/or the strength of each such relationship. Because such filtering operation are permutation-invariant logical operations, the graph neural network may therefore exhibit permutation-invariance of output based on the filtering operation for differently ordered but logically equivalent representations of the nodes (i.e., people) and edges (i.e., relationships) of the social network. In yet another example, some graph neural networks include an encoding or "bottleneck."

[3020] layer, in which an output from N neurons of a preceding layer is received as input and processed by a following layer that includes fewer than N neurons. Due to the smaller number of neurons in the following layer, the volume of data that encodes features of the output of the preceding layer is compressed into a smaller volume of data that encodes features of the output of the following layer. This compression of features, based on learned parameters and training of the

graph neural network to produce expected outputs, can cause the graph neural network to encode only more significant features of the processed data, and to discard less significant features of the processed data. The reduced-size output of the neurons of the following layer can be referred to as a latent space encoding of the input feature set. For example, whereas an input graph data set may include nodes that correspond to all pixels of an image of a cat, and an output of a previous layer of the graph neural network may include partially processed information about each node (i.e., each pixel) of the image of the cat, the output of the following layer of the graph neural network may include only features that correspond to visually significant features of the cat (e.g., features that correspond to the pixels that represent the distinctively shaped ears, eyes, nose, and mouth of the cat). Thus, the latent space encoding may reduce the processed input of the graph data set into a smaller encoding of nodes that represent significant visual features of the graph data set, and may exclude data about nodes that do not represent significant visual features of the graph data set. Many such graph neural networks include one or more "bottleneck" layers as one or more autoencoder layers, e.g., layers that automatically learn to generate latent space encodings of input data sets. As one such example, deep generative models may be used to generate output graph data that corresponds to various data types (e.g., images, text, video, scene graphs, or the like) based on an encoding, including an autoencoding, of an input such as a prompt or a random seed, Additional techniques for achieving or promoting permeation-invariance are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

[3021] In some scenarios, a graph data set may include a large number of nodes and/or a large number of edges. For example, a graph data set representing a social network may include thousands of nodes that represent people and millions of edges that represent relationships among the people. The size of the graph data set may result in an input vector that is very large (e.g., a very long input vector), and that might require a correspondingly large graph neural network to process (e.g., a graph neural network featuring millions of weights that connect the input graph data set to the nodes of a first layer of the graph neural network). The size of the input data set may result in large and perhaps prohibitive computational resources to receive and/or process the graph data set (e.g., large and costly storage and/or processing to store the input graph data set and/or the parameters and/or hyperparameters of the graph neural network, and/or a protracted delay in completing the processing of an input graph data set by the graph neural network). Further, the graph data set may exhibit properties of sparsity that cause a large portion of the input data set to be inconsequential. For example, a graph data set representing a social network may be encoded as a vector of N units respectively representing each node (i.e., each person) followed by a vector of NxN units that respectively represent a potential relationship between each node N1 and each node N2. Edges that represent a multidimensional mapping of connections between nodes (such as an NxN mapping of edges that represent possible connections between nodes) can be referred to as an adjacency matrix. However, in the social network, most person may have only a small number of relationships (i.e., far less than N-1 relationships with all other people of the social network). Thus, in the vector encoding of the input graph data set, a large majority of the NxN units that respectively represent potential relationships between each pair of nodes N1, N2 (i.e.,

the adjacency matrix) may be negative or empty (representing no relationship), and only a very small minority of the NxN units that respectively represent potential relationships between each pair of nodes N1, N2 may be positive or non-empty (representing a relationship). As another example, a graph data set representing a region may include N nodes representing cities and NxN edges representing possible roads between cities. However, if each city is only directly connected to a small number of neighboring cities, then a large majority of the NxN edges representing possible roads between cities (i.e., the adjacency matrix) may be negative or empty (representing no road connection), and only a very small minority of the NxN units that respectively represent potential roads between each pair of nodes N1, N2 may be positive or non-empty (representing an existing road). In such cases, the sparsity of an input vector representing the graph neural network may inefficiently consume computational resources (e.g., inefficiently applying storage and/or computation to large numbers of negative or empty units of the input vector) and/or may unproductively delay the completion of processing of the input graph data set.

[3022] Various techniques can be applied to reduce the sparsity of graph data sets and the processing of such graph data sets by graph neural networks. As a first example, the graph neural network can be pruned to reduce the number of nodes and/or edges included as an input data set (e.g., filtering the nodes of a graph neural network to a small cluster of densely related nodes, such as a small number of highly interrelated nodes that represent the members of a social circle in a social network). As a second example, the graph neural network can be encoded in a way that reduces sparsity. For example, rather than encoding the input graph data set as an adjacency matrix, the graph neural network may be configured to receive an encoding of the input graph data set as an adjacency list, i.e., as a list of edges that respectively connect two or more nodes of the graph. Due to encoding only information about existing edges, an adjacency list can eliminate or at least reduce the encoding of nonexistent edges. As a result, the size of the adjacency list may therefore be much smaller than the size of a corresponding adjacency matrix and can therefore eliminate or at least reduce the sparsity of the input graph data set. The adjacency list can include edge properties of the edges of the graph data set. The adjacency list can be limited to a particular size (e.g., the top N most influential connections in a social network). The nodes of the input graph data set can be limited based on the edges included in the adjacency list (e.g., excluding any nodes that are not connected to at least one of the edges included in the adjacency list). In yet another example, rather than encoding an entire set of nodes and edges, a graph neural network can be represented as an encoding of the nodes and edges. For example, a graph data set may include nodes that represent pixels of an image and edges that represent spatial representations of the pixels. However, if large areas of the image are inconsequential (e.g., dark, empty, or not associated with any notable objects in a segmented image), then large portions of the nodes and/or edges would be inconsequential. Instead, the image can be reencoded as a frequency-domain representation as coefficients associated with respective frequencies of visual features within the image. The frequency-domain representation may present greater information density than the adjacency matrix of pixels, and therefore may present an input to the graph neural network that encodes the visual features of the input graph data set with reduced sparsity.

[3023] Other techniques for eliminating or reducing sparsity, and therefore increasing efficiency, involve the architecture of the graph neural network. For example, the input graph data set may encode edges as an adjacency matrix, and a first layer of the graph neural network may reencode the edges of the input graph data set as an adjacency list for further processing by the graph neural network. As another example, the graph neural network may include a first one or more layers that is configured to process an entirety or at least a large portion of the nodes and/or edges of an input graph data set, followed by a filtering layer that is configured to limit an output of the first one or more layers of the graph neural network. For example, in a graph data set that includes nodes that represent people and edges that represent connections, a first one or more layers may process all of the nodes and/or edges, and a filtering layer can limit the further processing of the output of the first one or more layers to the nodes and/or edges for which the outputs of the first one or more layers are above a threshold (e.g., an influence and/or relationship significance threshold). As still another example, the graph neural network may receive a sparse graph input data set but may only process a portion of the input graph data set (e.g., one or more random sampling of subsets of nodes and/or edges). In some cases, the graph neural network may compare results of the processing of subsets of the input graph data set (e.g., randomly sampled subsets of the nodes and/or edges) and may aggregate such results until the results appear to converge within a confidence threshold. In this manner, the graph neural network may generate an acceptable output within the confidence threshold while avoiding processing an entirety of the sparse input graph data set. Many such techniques for eliminating and/or reducing sparsity are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

Input, Processing, and Output

[3024] Graph data sets may represent a variety of data types, including (without limitation) maps of geographic regions, including nodes representing cities and edges representing roads that connect two or more cities; social networks, including nodes representing people and edges representing relationships between two or more people; communication networks, including nodes representing people or devices and edges representing communication connections between the nodes or edges; economies, including nodes representing companies and edges representing transactions between two or more companies; molecules, including nodes representing atoms and edges representing bonds between two or more atoms; collections of events, including nodes representing individual events and edges representing causal relationships among two or more events; and periods of time, including nodes representing events and edges representing chronological periods among two or more events. Graph data sets may also be represent data types, such as passages of text, including nodes representing words and edges representing relationships among two or more words; images, including nodes representing pixels and edges representing spatial relationships among two or more pixels; object graphs, including nodes representing objects and edges representing dependencies among two or more objects; and three-dimensional spatial maps, including nodes representing three-dimensional objects and edges representing spatial relationships among two or more of the three-dimensional objects. Some graph data sets may include two or more subgraphs. In some such graph data sets, each node and/or each edge is

exclusively included in one subgraph. In some other graph data sets, at least one node and/or at least one edge may be included in two or more subgraphs, or in zero subgraphs. Some graph data sets are associated with non-graph data that is also included as input to a graph neural network. For example, a graph neural network that evaluates traffic patterns within a geographic region may receive, as input, both an input graph data set that includes nodes that represent cities and edges that represent roads interconnecting the cities, and also non-graph data representing traffic and/or weather features within the geographic region (e.g., traffic volume estimates and current or forecasted weather conditions that affect the traffic patterns).

[3025] As another example, some graph data sets may include an indication of zero or more cycles occurring among the nodes and/or edges of the graph data set. For example, a directed and/or undirected graph data set may include an indication that a particular cycle exists within the graph and includes a particular subset of nodes and/or edges. Alternatively, a directed and/or undirected graph data set may include an indication that the graph is acyclic and does not include any cycles. A graph neural network may be configured to receive, as input, and process a graph data set that includes an indication of zero or more cycles.

[3026] As another example, some graph data sets may include nodes for which the edges provide spatial dimensions. As a first example, in a graph representing a geographic region, nodes that represent cities are related by edges that represent distances, wherein the nodes and interrelated edges can form a spatial map of the geographic region. As a second example, in a graph representing a molecule, nodes that represent atoms are related by edges that represent chemical bonds between the atoms, and the arrangement of atoms by the bonds forms a three-dimensional molecular structure. In some such scenarios, the spatial relationships are well-defined by the nodes and edges. In other such scenarios, the spatial relationships can be inferred based on semantic relationships among the nodes and/or edges of the graph data set. For example, in a graph representing a language, nodes that represent words are related by edges that represent semantic relatedness of the words within a high-dimensional language space. A language model can generate an embedding of the words of the language in a multidimensional embedding space, wherein nodes that are close together within the embedding space represent synonyms, closely related concepts, or words that frequently appear together in certain contexts, whereas nodes that are not close together within the embedding space represent unrelated concepts or words that do not commonly appear together in various contexts. A variety of graph embedding models may be applied to this task, including (without limitation) DeepWalk, node2vec, line, and/or GraphSAGE. A graph neural network can be configured to receive, as input, an embedding of a graph data set instead of representations of the nodes and/or edges of the graph data set. Alternatively, a graph neural network can be configured to receive an input graph data set including representations of the nodes and/or edges of the graph data set, generate an embedding based on the input graph data set, and apply further processing to the embedding instead of to the input graph data set. A graph neural network that is configured to process an embedding instead of an input graph data set may exhibit greater permutation invariance (e.g., due to the semantic associations represented by the embedding) and/or increased efficiency due to reduced sparsity of the input.

[3027] Some graph data sets include representations of each of one or more nodes and each of one or more edges. Some graph neural networks are configured to receive and process such representations of graph neural networks. For example, the graph neural network may be configured to receive an input vector including an array of data representing each of the one or more nodes followed by an array of data representing each of the one or more edges, either as an adjacency matrix of possible edges between pairs of nodes or an adjacency list of existing edges. The input vector may encode the nodes and/or edges in a particular order (e.g., a priority order of nodes and/or a weight order of edges) or in an unordered manner. Alternatively or additionally, the graph data set may include and/or encode other types of information about each of one or more nodes and/or each of one or more edges of the graph data set. For example, the graph may include a hierarchical organization of nodes and/or edges relative to one another and/or to a fixed reference point. The graph neural network may be configured to receive and process an input graph data set that includes an indication of the arrangement of one or more nodes and/or one or more edges in the hierarchical organization.

[3028] As another example, a graph may include an indication of a centrality of one or more nodes and/or edges within the graph (e.g., a graph of a social network including nodes that are ranked based on a centrality of each node to a cluster). The graph neural network may be configured to receive and process an input graph data set that includes an indication of a centrality of one or more nodes and/or one or more edges in the graph.

[3029] As another example, a graph may include an indication of a degree of connectivity of one or more nodes and/or edges within the graph (e.g., a graph of a social network including nodes that are ranked according to a count of other nodes to which each node is connected by one or more edges, and/or a degree of significance of a relationship represented by an edge based on the nature of the relationship and/or the degrees of the nodes connected by the edge). The graph neural network may be configured to receive and process an input graph data set that includes an indication of a degree of one or more nodes and/or one or more edges in the graph.

[3030] As another example, a graph may include an indication of one or more clusters occurring within the graph. For example, a graph may include a result of a clustering analysis of the graph, e.g., a determination of k clusters within the graph and an identification of the nodes and/or edges that are included in each cluster. The clusters may be determined by a k-means clustering analysis, a Gaussian mixture model of with variable numbers of clusters and variable Gaussian orders, or the like. A graph may include a clustering coefficient of one or more nodes and/or one or more edges (e.g., a measurement of a degree to which at least some of the nodes and/or edges of a subgraph of the graph are clustered based on similarity and/or activity). The graph neural network may be configured to receive and process an input graph data set that includes an indication of a clustering coefficient of one or more nodes and/or one or more edges in the graph or a subgraph thereof.

[3031] As another example, a graph may include an indication of a graphlet degree vector that indicates a graphlet that is represented one or more times in the graph. For example, in a graph representing atoms in a regular structure such as a crystal, the graph may include a graphlet degree

vector that indicates and/or describes a graphlet representing a recurring atomic structure, and an encoding of the regular structure that indicates each of one or more occurrences of a graphlet, including a location and/or orientation, and/or a count of occurrences of the graphlet. The graph neural network may be configured to receive and process an input graph data set that includes a graphlet degree vector, and, optionally, features of one or more occurrences of a graphlet in the graph and/or a count of the occurrences of the graphlet in the graph.

[3032] As another example, a graph may include an indication of one or more paths and/or traversals of one or more nodes and/or one or more edges of the graph, optionally including additional details associated with a path or traversal such as a popularity, frequency, length, difficulty, cost, or the like. For example, in a graph representing a spatial arrangement of nodes, the graph may include a path or traversal of edges that connect a first node to a second node through zero or more other nodes, as well as properties of the path or traversal such as a total length, distance, time, and/or cost. The graph neural network may be configured to receive and process an input graph data set that includes additional details associated with one or more paths or traversals, including an indication (e.g., a list) of the associated nodes and/or edges and a list of one or more properties of the path and/or traversal.

[3033] As another example, a graph may include an indication of metrics or properties that relate one or more nodes and/or one or more edges. For example, in a graph including a spatial arrangement of nodes, the graph may include an indication of a shortest distance between two nodes and/or an indication of a set of nodes and/or edges that are common to two nodes. As another example, a graph representing a network of communicating devices may include a routing table of one or more routes that respectively indicate, for a particular node and a particular edge connected to the node, a list of other nodes and/or edges that can be efficiently reached by traversing based on the particular edge. As yet another example, in a graph representing a social network including nodes that represent people, the graph may indicate, for at least one pair of nodes, a measurement of similarity of the nodes based on their node properties, edges, locations in the social network, connections to other nodes, or the like (e.g., a Katz index of node similarity) and/or, for at least one pair of edges, a measurement of similarity of the edges based on their edge properties, connected nodes, locations in the social network, or the like (e.g., a Katz index of edge similarity). The graph neural network may be configured to receive and process an input graph data set that includes one or more metrics or properties that relate one or more nodes and/or one or more edges (e.g., a routing table of routes within the graph, and/or a Katz index that indicates a measurement of similarity among at least two nodes and/or at least two edges).

[3034] As another example, a graph may include an indication of various graph properties of the graph (e.g., a graph size, graph density, graph interconnectivity, graph chronological period, graph classification, a count of subgraphs within the graph, or the like). For example, in a graph including two or more subgraphs (e.g., a social network including two or more social circles), the graph data set may include a measurement of a similarity of each subset of at least two subgraphs of the graph. The measurement of the similarity may be determined based on one or more graph kernel methods (e.g., a Gaussian radial basis function that can be applied to the graph to identify one or more

clusters of similar nodes that comprise a subgraph). As another example, a graph may include a measurement of similarity with respect to another graph (e.g., an indication of whether a particular social network graph resembles other social network graphs that have been classified as representing a genealogy or lineage, a set of friendships, and/or a set of professional relationships). The graph neural network may be configured to receive and process an input graph data set that includes measurements determined by one or more graph properties (e.g., one or more measurements of similarity of one or more nodes, edges, and/or subgraphs, and/or a measurement of similarity of the graph to other graphs). Further explanation and/or examples of various graph data sets that may be provided as input to graph neural networks are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

[3035] Graph neural networks may be configured to perform various types of processing over such graph data sets. As previously discussed, a graph neural network can be organized as a series of layers, each of which can include one or more nodes that receive input, apply an activation function, and generate output. The output of each node of a first layer can be multiplied by a weight of a connection between the node and a node of a second layer, and then added to a bias associated with the first layer, to generate an input to the node of the second layer. The graph neural network can include various additional layers that perform other types of processing, including (without limitation) pooling, filtering, and/or latent space encoding operations, memory or stateful features, and recurrent and/or reentrant processing.

[3036] Some graph neural networks may perform label propagation among the nodes and/or edges of a graph data set. For example, in an input graph data set, one or more nodes and/or one or more edges may be associated with one or more labels of a label set, while one or more other nodes and/or one or more other edges may not be associated with any labels. A graph neural network may apply a label propagation algorithm (LPA) to assign labels to one or more unlabeled nodes and/or one or more unlabeled edges. For example, the graph neural network may assign a label to an unlabeled node based on labels associated with one or more edges connected to the node, and/or with one or more other nodes that are connected to the node by the one or more edges. The graph neural network may assign a label to an unlabeled edge based on labels associated with one or more nodes connected by the edge, and/or with one or more other edges that are also connected to the nodes connected by the edge. Some graph neural networks may perform label propagation based on a voting, consensus, weighting, and/or scoring determination. For example, a graph neural network may be unable to perform a classification of an unlabeled node and/or unlabeled edge based solely on the node properties and/or edge properties, but may be able to perform the classification based on a further consideration of the labels associated with other nodes and/or edges within a neighborhood of the unlabeled node and/or unlabeled edge.

[3037] Some graph neural sets may perform a scoring and/or ranking of nodes and/or edges of a graph data set. As an example, in a graph data set that represents the World Wide Web and that includes nodes that represent web pages and directed edges that represent hyperlinks of linking web pages to linked web pages, a graph neural network may determine one or more scores of each node (i.e., each web page) based on the scores of other nodes that hyperlink to the node. Each score

may further be based on the scores of the other nodes that include a directed edge to this node (e.g., the scores of other web pages that hyperlink to this page). Additionally, each score associated with a node may represent a weight of an association between the web page and a particular topic (e.g., a particular topic or keyword that is associated with the web page, hyperlinks, and/or other pages that hyperlink to this web page). In some cases, the scores may be personalized based on the activities of a particular user (e.g., based on the hyperlinks from pages that the user frequently visits). A search engine may use the scores as rankings in order to generate search results for web searches including various topics or keywords (e.g., in response to a web search for a particular search term, present search results that correspond to the nodes with the highest scores associated with the search term, and present the search results in ranked order based on the scores). As another example, for a graph data set representing a social network, a graph neural network may generate a reputation score for each node based on other nodes that are associated with the node and the reputation scores of such other nodes. The scores of the nodes may be used to recommend new connections in the social network (e.g., recommending a first person connect with a second person, based on a high reputation score of the second person by people who are closely associated with the first person).

[3038] Some graph neural networks may perform a clustering analysis of the nodes and/or edges of a graph data set. As a first example, in a graph data set representing a social network, a graph neural network may perform a clustering analysis of the nodes representing the people of the social network, based on edges representing relationships among two or more nodes, in order to identify one or more clusters that represent social circles of highly interconnected people within the social network. Based on this clustering analysis, the graph neural network may partition the social network into subgraphs that respectively represent social circles, and may perform further, finer-grained evaluation of each social circle and the people represented by the nodes in each subgraph. As a second example, in a graph data set representing a social network, a graph neural network may perform a clustering analysis of the edges representing the relationships among people of the social network, in order to identify one or more clusters that represent different types of relationships, such as familial relationships, friendships, and professional relationships. Based on this clustering analysis, the graph neural network may partition the social network into subgraphs that respectively represent different types of social networks and may perform further analysis of relationships among two or more individuals based on the type of relationship associated with the subgraph to which the relationship belongs. In these and other scenarios, in order to perform clustering analysis, a graph neural network may utilize a variety of clustering algorithms. As one such example, a graph neural network may apply spectral clustering techniques, wherein a similarity matrix that represents similarities among nodes and/or edges is evaluated to identify eigenvalues that indicate significant similarity relationships. Based on the similarity matrix, the graph neural network may perform a dimensionality reduction of the graph data set (e.g., reducing the features of the nodes and/or edges that are evaluated to determine clusters in order to focus on features that are highly correlated with and/or indicative of significant similarities). Dimensionality reduction of the graph data set based on the similarity matrix may enable the graph neural network

to determine clusters more efficiently and/or rapidly, e.g., by reducing a high-dimensionality graph data set (wherein each node and/or edge is characterized by a multitude of node properties and/or edge properties) into a lower-dimensionality graph data set of a subset of features that are highly correlated with and/or indicative of similarity and clustering.

[3039] Some graph neural networks may perform a centrality determination among nodes and/or edges of a graph data set. For example, for a graph data set representing a social network, a graph neural network may evaluate the graph data set to identify a subset of nodes based on a centrality among the edges representing the connections of the social network, e.g., people who are at the center of each of one or more social circles within the social network. Alternatively or additionally, some graph neural networks may perform a "betweenness" determination among the nodes and/or edges of the graph data set. For example, a node may be considered to be "between" two clusters of nodes, such as a member of two or more clusters representing two or more social circles. Such "between" nodes may represent a communication bridge that conducts information between clusters (e.g., a person who can convey ideas and/or influence from a first social circle to a second social circle and vice versa). Some such graph neural networks may perform "betweenness" determinations based on a betweenness centrality measurement, e.g., based on a measurement of a shortest path between all pairs of nodes in the graph data set. As another example, a graph data set may represent a collection of text documents, wherein each node represents a document and each edge represents a relationship between documents (e.g., a unidirectional or bidirectional citation between a first document and a second document). A graph neural network can perform a centrality determination and/or a betweenness determination to determine significant documents within the collection (e.g., a document that is heavily cited by one or more clusters of other documents, and/or a document that includes ideas or associations between the documents of a first cluster and the documents of a second cluster).

[3040] Some graph neural networks may perform analyses of structures occurring within a graph neural network. As an example, for a graph data set that represents a social network, a graph neural network may determine a notable sequence of relationships, such as a first relationship between node N1 and node N2 based on a shared interest, a second relationship between node N2 and node N3 based on the same shared interest, and a third relationship between node N3 and node N4 based on the same shared interest. Based on this sequence or chain of relationships, the graph neural network may recommend to a person represented by node N1 some further relationships with the people represented by nodes N3 and N4, due to the combination of shared interests and mutual relationships. In some such cases, a graph neural network may perform such structural analysis based on a traversal algorithm that traverses a sequence of nodes connected by one or more edges, and/or that traverses a sequence of edges connected by one or more nodes. As an example, a graph neural network may perform a random walk within the graph data set, such as starting with a first node (e.g., a first person of a social network) and following a limited set of edges that connect the first node to other nodes. In some cases, the traversal may be random (e.g., traversing from a node based on a random selection among the edges that connect the node to other nodes). In some other cases, the traversal may be weighted (e.g., each edge may include an edge property including a

weight that represents a strength of a relationship among two or more nodes, and the traversal may be based on a weighted random selection that preferentially selects higher-weighted connections over lower-weighted connections). In some cases, the traversal can include a restart probability, e.g., a probability of retrying the traversal beginning with the original node or another node, based on a score such as a distance of the traversal with respect to the original node. In these and other cases, the results of a random walk can be used in further analyses and/or activities of the graph neural network (e.g., presenting recommendations for new social connections among the nodes of a social network).

[3041] Some graph neural networks may perform an analysis of a graph data set based on an attention model. For example, in a social network, the influence of a particular person P1 may not be determined by the connectedness of person P1 to other people in the social network, but based on a perception of person P1 by other people of the social network as being knowledgeable, skilled, influential, or the like. Thus, a graph neural network may be configured to evaluate a graph data set representing a social network in which nodes represent people and edges represent relationships, but may be unable to determine influence based only on graph concepts such as connectedness of the nodes based on the edges. Rather, the graph neural network might model influence as an attention of each node (i.e., a second person P2 of the social network) upon each other node (e.g., person P1 of the social network). Thus, a particular opinion of person P2 of the social network may depend not only on the connections of person P2 to other people of the social network (including person P1), but also upon the attention that person P2 accords to such other people of the social network (including person P1). That is, even though person P2 is closely connected to certain people of the social network by various edges, the opinion of person P2 may be heavily shaped by person P1 and other people to whom person P2 is only indirectly connected in the social network. As a second such example, in a graph data set that represents traffic flow within a region, an edge E1 (e.g., a first road) may be directly connected to other edges of the graph data set, but an edge property of the edge E1 (e.g., a traffic volume and/or congestion of the road) may be impacted more heavily by edge properties of other edges to which edge E1 is not directly connected (e.g., roads in other parts of the geographic region for which traffic volume and/or congestion is highly determinative of the traffic volume and/or congestion of this road). Thus, in order to predict and/or estimate a traffic volume and/or congestion of a particular road, a graph neural network may evaluate not only the traffic volume and/or congestion of other roads that are directly connected to the particular road, but also other roads for which traffic volume and/or congestion is highly determinative of corresponding conditions of this road. In these and other scenarios, a graph neural network may evaluate a graph data set based on an attention model, in which analyses and updates of the state of nodes and/or edges of the graph data set are based, at least in part, on an attention of each node and/or edge upon other nodes and/or edges of the graph data set. For example, the graph neural network may include an attention layer that determines, for a particular node and/or edge of an input graph data set, which other nodes and/or edges of the input graph data set are likely to be relevant to determining an updated state of the particular node and/or edge. Various attention models may be used by such graph neural networks, including

multi-head attention models in which each node and/or edge is related to a plurality of other nodes and/or other edges with varying weighted attention values (e.g., by each of a plurality of attention layers). Multi-head attention models can allow a graph neural network to consider the influences upon a particular node and/or edge of a plurality of other nodes and/or edges, which may (or may not) be further related to one another by the graph structure and/or attention. Based on the attention model and the attention layers included in the graph neural network, the graph neural network can perform a more sophisticated graph analysis that is based on more than the structural relationships of the graph.

[3042] Some graph neural networks may be configured to process a graph data set in order to determine, and optionally output, various types of data (e.g., measurements, calculations, inferences, explanations, or the like) that relate to one or more nodes, one or more edges, and/or one or more subgraphs of the input graph data set and/or to the input data graph set as a whole. Some graph neural networks are configured to generate, and optionally output, various types of representations of graph neural networks. For example, the graph neural network may be configured to determine, and optionally output, an output vector including an array of data representing each of the one or more nodes followed by an array of data representing each of the one or more edges, either as an adjacency matrix of possible edges between pairs of nodes or an adjacency list of existing edges. The output vector may encode the nodes and/or edges in a particular order (e.g., a priority order of nodes and/or a weight order of edges, or corresponding to a corresponding order of the nodes and/or edges in the input graph data set) or in an unordered manner. Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, other types of information about each of one or more nodes and/or each of one or more edges of the graph data set. For example, the graph neural network may be configured to determine, and optionally output, a hierarchical organization of nodes and/or edges relative to one another and/or to a fixed reference point. Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, an output graph data set that includes an indication of the arrangement of one or more nodes and/or one or more edges in the hierarchical organization.

[3043] As another example, a graph neural network may be configured to determine, and optionally output, an indication of a centrality of one or more nodes and/or edges within the input graph data set (e.g., a graph of a social network including nodes that are ranked based on a centrality of each node to a cluster). Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, an output graph data set that includes an indication of a centrality of one or more nodes and/or one or more edges in the graph.

[3044] As another example, a graph neural network may be configured to determine, and optionally output, an indication of a degree of connectivity of one or more nodes and/or edges of an input graph data set (e.g., a graph of a social network including nodes that are ranked according to a count of other nodes to which each node is connected by one or more edges, and/or a degree of significance of a relationship represented by an edge based on the nature of the relationship and/or the degrees of the nodes connected by the edge). Alternatively or additionally, the graph

neural network may be configured to determine, and optionally output, an output graph data set that includes an indication of a degree of one or more nodes and/or one or more edges in the output graph data set.

[3045] As another example, a graph neural network may be configured to detect, identify, and/or analyze one or more clusters occurring within an input graph data set. For example, a graph neural network may be configured to perform a clustering analysis of an input graph data set to determine, and optionally output, a determination of k clusters within the input graph data set and an identification of the nodes and/or edges that are included in each cluster. The graph neural network may be configured to determine clusters based on a k-means clustering analysis, a Gaussian mixture model of with variable numbers of clusters and variable Gaussian orders, or the like. The graph neural network may be configured to determine, and optionally output, an indication of a clustering coefficient of one or more nodes and/or one or more edges of an input graph data set (e.g., a measurement of a degree to which at least some of the nodes and/or edges of a subgraph of the graph are clustered based on similarity and/or activity). Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, an output graph data set that includes an indication of one or more clusters including one or more nodes and/or one or more edges in the output graph data set or a subgraph thereof (e.g., a result of a k¬¨-means clustering analysis of an output graph data set, a Gaussian mixture model of an output graph data set, and/or one or more clustering coefficients of an output graph data set).

[3046] As another example, a graph neural network may be configured to determine, and optionally output, an indication of a graphlet degree vector that indicates a graphlet that is represented one or more times in an input graph data set. For example, for a graph representing atoms in a regular structure such as a crystal, the graph neural network may be configured to determine, and optionally output, a graphlet degree vector that indicates and/or describes a graphlet representing a recurring atomic structure, and an encoding of the regular structure that indicates each of one or more occurrences of a graphlet, including a location and/or orientation, and/or a count of occurrences of the graphlet in the input graph data set. Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, an output graph data set that includes a graphlet degree vector, and, optionally, features of one or more occurrences of a graphlet in the output graph data set and/or a count of the occurrences of the graphlet in the output graph data set.

[3047] As another example, a graph neural network may be configured to determine, and optionally output, an indication of one or more paths and/or traversals of one or more nodes and/or one or more edges of the input graph data set, optionally including additional details associated with a path or traversal such as a popularity, frequency, length, difficulty, cost, or the like. For example, for an input graph data set representing a spatial arrangement of nodes, the graph neural network may be configured to determine, and optionally output, a path or traversal of edges that connect a first node to a second node through zero or more other nodes of the input graph data set, as well as properties of the path or traversal such as a total length, distance, time, and/or cost. Alternatively or additionally, the graph neural network may be configured to determine, and

optionally output, an output graph data set that includes additional details associated with one or more paths or traversals, including an indication (e.g., a list) of the associated nodes and/or edges of the output graph data set and a list of one or more properties of each such path and/or traversal.

[3048] As another example, a graph neural network may be configured to determine, and optionally output, an indication of metrics or properties that relate one or more nodes and/or one or more edges of an input graph data set. For example, for an input graph data set including a spatial arrangement of nodes, the graph neural network may be configured to determine, and optionally output, an indication of a shortest distance between two nodes and/or an indication of a set of nodes and/or edges that are common to two nodes of the input graph data set. As another example, for an input graph data set representing a network of communicating devices, the graph neural network may be configured to determine, and optionally output, a routing table of one or more routes that respectively indicate, for a particular node of the input graph data set and a particular edge connected to the node, a list of other nodes and/or edges of the input graph data set that can be efficiently reached by traversing based on the particular edge. As yet another example, for an input graph data set representing a social network including nodes that represent people, the graph neural network may be configured to determine, and optionally output, an indication for at least one pair of nodes of a measurement of similarity of the nodes of the input graph data set based on their node properties, edges, locations in the social network, connections to other nodes, or the like (e.g., a Katz index of node similarity) and/or, for at least one pair of edges of the input graph data set, a measurement of similarity of the edges based on their edge properties, connected nodes, locations in the social network, or the like (e.g., a Katz index of edge similarity). Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, an output graph data set that includes one or more metrics or properties that relate one or more nodes and/or one or more edges (e.g., a routing table of routes within the graph, and/or a Katz index that indicates a measurement of similarity among at least two nodes and/or at least two edges).

[3049] As another example, a graph neural network may be configured to determine, and optionally output, an indication of various graph properties of an input graph data set (e.g., a graph size, graph density, graph interconnectivity, graph chronological period, graph classification, a count of subgraphs within the graph, or the like). For example, for an input graph data set including two or more subgraphs (e.g., a social network including two or more social circles), the graph neural network may be configured to determine, and optionally output, a measurement of a similarity of each subset of at least two subgraphs of the input graph data set. The measurement of the similarity may be determined based on one or more graph kernel methods (e.g., a Gaussian radial basis function that can be applied to the input graph data set to identify one or more clusters of similar nodes that comprise a subgraph). As another example, a graph neural network may be configured to determine, and optionally output, a measurement of similarity of an input graph data set with respect to another graph data set (e.g., an indication of whether a particular social network graph resembles other social network graphs that have been classified as representing a genealogy or lineage, a set of friendships, and/or a set of professional relationships). Alternatively or additionally, the graph neural network may be configured to determine, and optionally output, an

output graph data set that includes measurements determined by one or more graph properties of an output graph data set (e.g., one or more measurements of similarity of one or more nodes, edges, and/or subgraphs, and/or a measurement of similarity of the output graph data set to the input graph data set and/or other graph data sets). Further explanation and/or examples of various types of processing that graph neural networks can determine, and optionally output, for various input graph data sets and/or output graph data sets are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

[3050] Graph neural networks may be configured to generate various forms of output that correspond to various tasks. For example, graph neural networks can generate output that represents node-level predictions that relate to one or more nodes of an input graph data set. The node-level predictions can include a discovery of a new node that was not included in the input graph data set. For example, in a graph data set including edges that represent travel of individuals in a region, the nodes can represent points of interest, and the graph neural network can discover a new node that corresponds to a new point of interest. The node-level predictions can include an exclusion of a node that is included in the input graph data set. For example, in a graph data set including edges that represent travel of individuals in a region, the nodes can represent points of interest, and the graph neural network can exclude an existing node that no longer represents a point of interest. The node-level predictions can include a classification of a node that is included in the input graph data set, or of a newly discovered node that was not included in the input graph data set (e.g., a classification of the node as being of a node type selected from a set of node types, as being associated with one or more labels of a classification label set, and/or as belonging to zero or more subgraphs of the graph data set). For example, in a graph data set representing locations within a geographic region, the graph neural network can generate a prediction of a classification of a location of interest as one or more particular types of locations of interest (e.g., a source of food, a source of fuel, a lodging location, and/or a tourist destination). The node-level predictions can include an identification of a node from among the nodes of the input graph data set based on various features, or of a newly discovered node. For example, in a graph data set representing a social network and including nodes that represent people, the graph neural network can identify a particular node that corresponds to a particular person, such as an influential person of the social network. The node-level predictions can include a determination and/or updating of one or more node properties of one or more existing and/or newly discovered nodes, such as a prediction of a demographic feature, opinion, or interest of a node representing a person in a social network.

[3051] As another example, graph neural networks can generate output that represents edge-level predictions that relate to one or more edges of an input graph data set. The edge-level predictions can include a discovery of a new edge that was not included in the input graph data set. For example, in a graph data set representing a social network that includes nodes that represent people, a graph neural network can output a prediction (e.g., a recommendation) of a relationship between two nodes that correspond to two people in a small social circle of highly interconnected people. The node-level predictions can include an exclusion of a node that is included in the input graph data set. For example, in a graph data set representing a social network that includes nodes that

represent people, a graph neural network can output a prediction of a no-longer-existing edge that corresponds to a relationship that no longer exists (e.g., a lost connection based on a splitting of a social circle). The edge-level predictions can include a classification of an edge that is included in the input graph data set, or of a newly discovered edge that was not included in the input graph data set (e.g., a classification of an edge as being a of an edge type selected from a set of edge types, as being associated with one or more labels of a classification label set, and/or as belonging to zero or more subgraphs of the graph data set). For example, in a graph data set representing a social network, a graph neural network can generate a predicted classification of an edge as representing a relationship between two people as of one or more relationship types (e.g., a familial relationship, a friendship, or a professional relationship). The edge-level predictions can include an identification of an edge from among the edges of the input graph data set based on various features, or of a newly discovered edge. For example, in a graph data set representing a social network and including edges that represent relationships, the graph neural network can identify a particular edge that corresponds to a potential relationship to be recommended to the associated people, such as two people of the social network who are not yet connected but who share common personal or professional interests. The edge-level predictions can include a determination and/or updating of one or more edge properties of one or more existing and/or newly discovered edges, such as a prediction of a demographic feature, opinion, or interest that serves as the basis for a relationship between two people of the social network.

[3052] As another example, graph neural networks can generate output that represents graph-level predictions that relate to one or more graph properties of the input graph data set. The graph-level predictions can include a discovery of a new graph property that was not associated with the input graph data set. For example, in a graph data set representing a social network that includes nodes that represent people and edges that represent relationships, a graph neural network can output a prediction of a demographic trait, opinion, or interest that is common or popular among the people of the social network, or a relationship behavior that is exhibited in the relationships among the people of the social network. The graph-level predictions can include an exclusion of a graph property that was associated with the input graph data set. For example, in a graph data set representing a social network that includes a graph property based on a shared interest, a graph neural network can output a prediction that the interest no longer appears to be common and/or popular among the people of the social network, or of a relationship behavior that is no longer exhibited among the relationships of the people of the social network. The graph-level predictions can include a classification of the input graph data set (e.g., a classification of the graph data set, or at least a portion thereof, as being associated with one or more labels of a classification label set). For example, in a graph data set representing a social network, a graph neural network can generate a predicted classification of the graph as representing a familial social network, a friendship social network, and/or a professional social network. The graph-level predictions can include an identification of one or more subgraphs of the graph based on common features of the nodes and/or edges included in the subgraph. For example, in a graph data set representing a social network, the graph neural network can subgraphs that correspond to various social circles of highly

interconnected people. The graph-level predictions can include a determination and/or updating of one or more graph properties of the graph, such as an updating of a frequency of communication and/or a strength of relationships among the people of a social network.

[3053] As another example, graph neural networks can perform graph-to-graph translation by receiving an input graph data set and generating output that represents a different graph data set. For example, a graph neural network can receive an input graph data set and can generate an output graph data set that includes one or more newly discovered nodes and/or edges; an exclusion of one or more nodes and/or edges; a classification of one or more nodes and/or edges; an identification of one or more nodes and/or edges; and/or an update of one or more node properties, edge properties, and/or graph properties. A graph neural network can receive an input graph data set and can generate an output graph data set that shares various similarities with the input graph data set. For example, a graph neural network can receive, as input, a first graph representing a first geographic region (e.g., a real geographic region) and can generate, as output, a first graph representing a different geographic region (e.g., a fictitious geographic region) that shares similarities with the first graph and that has some dissimilarities with respect to the first graph. A graph neural network can receive, as input, an input graph data set and can generate, as output, a subgraph of the input graph data set. A graph neural network can receive, as input, an input graph data set and can generate, as output, an expanded graph including a first subgraph corresponding to the input graph data set and a second subgraph that is newly generated. A graph neural network can receive, as input, a first graph that corresponds to a first time and can generate, as output, a second graph that corresponds to a different time than the first time. For example, the graph neural network can receive, as input, a graph data set that corresponds to a state of a geographic region at a current time, and can generate, as output, a graph data set that predicts the state of the geographic region at a past time or a future time.

[3054] As another example, graph neural networks can generate graphs from non-graph input data. For example, a graph neural network can receive, as input, locations of travelers within a geographic region over a period of time, and can generate, as output, graph data that includes one or more nodes that represent points of interest among the travelers and edges that represent paths between the points of interest (e.g., roads that connect the points of interest). As another example, a graph neural network can receive, as input, a description of a graph (e.g., a natural-language description of a geographic location) and can generate, as output, graph data that corresponds to the description of the graph (e.g., a graph of a region that includes one or more nodes representing locations and one or more edges representing roads that interconnect the locations). The graph neural network may receive both graph data and non-graph data (e.g., a graph representing a social network and an indication of a particular person in the social network) and can generate, as output, graph data based on the input (e.g., a subgraph of the people who consider the identified person to be influential).

[3055] As another example, graph neural networks can receive an input graph data set and can generate, as output, non-graph data. For example, a graph neural network can receive, as input, a graph representing a social network including nodes that represent people and edges that represent

relationships, and can generate, as output, one or more metrics of the social network (e.g., an average number of connections among the people of the social network, an identification of a person of high influence within the social network, or a description of a relationship behavior that commonly occurs within the social network). As another example, a graph neural network can receive, as input, a graph representing a geographic region including nodes that represent locations and edges that represent roads connecting the locations, and can generate, as output, one or more predictions and/or measurements of traffic within the geographic region. The graph neural network may receive both graph data and non-graph data (e.g., a graph representing a social network and an indication of a particular person in the social network) and can generate, as output, non-graph data based on the input (e.g., a summary and/or prediction of the social behaviors of the identified person). For example, a graph neural network that evaluates traffic patterns within a geographic region may process, and optionally output, both an output graph data set that includes nodes that represent cities and edges that represent roads interconnecting the cities, and also non-graph output data representing predictions and/or inferences of traffic and/or weather features within the geographic region (e.g., traffic volume estimates and current or forecasted weather conditions that affect the traffic patterns).

[3056] As another example, some graph neural networks may be configured to determine, and optionally output, an indication of zero or more cycles occurring among the nodes and/or edges of an input graph data set. For example, for a directed and/or undirected input graph data set, a graph neural network may determine, and optionally output, an indication that a particular cycle exists within the input graph data set and includes a particular subset of nodes and/or edges. Alternatively, for a directed and/or undirected graph data set, a graph neural network may determine, and optionally output, an indication that the graph is acyclic and does not include any cycles. A graph neural network may be configured to determine, and optionally output, an output graph data set that includes an indication of zero or more cycles.

[3057] As another example, graph neural networks can receive an input graph data set and can generate, as output, an interpretation and/or explanation of the input graph data set. For example, a graph neural network can receive, as input, a graph representing a collection of devices, including nodes that respectively represent a device and edges that respectively represent an instance of communication and/or interaction among two or more devices. The graph neural network can generate, as output, an interpretation and/or explanation of the communications and/or interactions represented in the graph, such as an explanation of a set of interactions as being part of a collective and/or collaborative effort among the two or more devices and/or a related series of interactions that are associated with a particular activity. The explanation and/or interpretation may include, for example, a classification of one or more nodes, edges, patterns of activity, and/or the graph; a natural-language summary or narrative explanation of one or more nodes, edges, patterns of activity, and/or the graph; a data set that characterizes one or more nodes, edges, patterns of activity, and/or the graph; and/or a presentation (e.g., a static or motion visualization) of one or more nodes, edges, patterns of activity, and/or the graph. As one such example, a graph neural network may identify, within an input graph data set, one or more subgraphs (e.g., one or more

clusters of related nodes and/or edges), and may output an interpretation and/or explanation of the subgraph (e.g., a description of the set of features that characterize the subgraph or cluster). As another example, a graph neural network may generate a visualization of a subgraph of an input graph data set, wherein the visualization depicts, highlights, and/or illustrates a structure and/or an anomalous feature of the subgraph. Some such graph neural networks may be configured to generate interpretations and/or explanations of any input graph data set, e.g., based on an identification of features of an input data set that inform such interpretations and/or explanations, such as clusters, outliers, or determinations of apparent structure and/or data relationships. Other such graph neural networks may be configured to generate domain-specific interpretations and/or explanations of domain-specific graph data sets. For example, a graph neural network may be configured to analyze a graph data set representing a social network identify both a subset of the social network corresponding to an influential cluster of people of the social network and also an interpretation and/or explanation of why this cluster of people appears to be influential within the social network. Graph neural networks can generate interpretations and/or explanations using a variety of techniques, including "white-box" analysis techniques that can be applied to various properties of graph data sets and components thereof. Examples of graph neural networks that include instance-level explanations based on gradients and/or features include, without limitation, Guided BP, class activation mapping (CAM), and GradCAM. Examples of graph neural networks that include instance-level explanations based on perturbations include, without limitation, GNNExplainer, PGExplainer, ZORRO, and Graphmask. Examples of graph neural networks that include instance-level explanations based on decomposition include, without limitation, layer-wise relevance propagation (LRP), Excitation BP, and GNN LRP. Examples of graph neural networks that include instance-level explanations based on surrogate analysis include, without limitation, GraphLIME, RelEX, and PGMExplainer. Examples of graph neural networks that include model-level explanations include XGNN. Further explanation and/or examples of various interpretable and/or explainable features of graph data sets or components thereof that may be generated by graph neural networks are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

Architectures and Frameworks

[3058] Graph neural networks may be designed and/or organized according to various architectures. For example, a multilayer graph neural network may include a number of layers, each layer including a number of neurons. In each layer of the graph neural network, the neurons may be configured to receive, as input, at least a portion of an input data set (e.g., an input graph data set) and/or at least a portion of an output of at least one neuron of one or more layers of the graph neural network. Additionally, in each layer of the graph neural network, the neurons may be configured to generate, as output, at least a portion of an output data set of the graph neural network (e.g., an output graph data set of graph neural network) and/or at least a portion of an input to at least one neuron of one or more layers of the graph neural network.

[3059] In some graph neural networks, an architecture of the graph neural network is based on the input to the graph neural network. For example, a fixed-size graph of N nodes and E edges

interconnecting the nodes may be received and processed by a graph neural network that includes an input layer featuring N neurons respectively configured to receive input from one of the N nodes and/or E neurons respectively configured to receive input from one of the E edges. A graph including an adjacency list having a maximum of E edges may be received and processed by a graph neural network that includes an input layer featuring E neurons respectively configured to receive and process one of the E edges represented in the adjacency list. A graph including two subgraphs may be received and processed by a graph neural network that includes an input layer featuring a first set of neurons that are configured to process the nodes and/or edges of the first subgraph and a second set of neurons that are configured to process the nodes and/or edges of the second subgraph. In some graph neural networks, an architecture of the graph neural network may be based on non-graph input data that is received and processed by the graph neural network. For example, a graph neural network may be configured to receive, as input, a description of a graph (e.g., a number of nodes and/or edges and one or more properties of the graph). The graph neural network may be further configured to generate a graph corresponding to the description, and to process and optionally output the graph according to various graph neural network processing techniques.

[3060] In some graph neural networks, an architecture of the graph neural network is based on an output of the graph neural network. For example, a graph neural network may be configured to determine, and optionally output, a fixed-size output graph data set including N nodes and E edges. The graph neural network may therefore include an output layer featuring N neurons respectively configured to generate output corresponding to one of the N nodes and/or E neurons respectively configured to generate output corresponding to one of the E edges. A graph neural network may be configured to determine, and optionally output, an adjacency list having a maximum of E edges. The graph neural network may therefore include an output layer featuring E neurons that respectively generate output corresponding to one of the E edges represented in the adjacency list. A graph neural network may be configured to determine, and optionally output, an output graph data set including two subgraphs. The graph neural network may therefore include an output layer featuring a first set of neurons that are configured to generate output corresponding to the nodes and/or edges of the first subgraph and a second set of neurons that are configured to process the nodes and/or edges of the second subgraph. In some graph neural networks, an architecture of the graph neural network may be based on non-graph output data that is determined, and optionally output, by the graph neural network. For example, a graph neural network may be configured to determine, and optionally output, a description of an input graph data set and/or an output graph data set (e.g., a number of nodes and/or edges and one or more properties of the input graph data set and/or the output graph data set), according to various graph neural network processing techniques.

[3061] In some graph neural networks, an architecture of the graph neural network may be based on a directionality of one or more edges included in an input data set and/or an output data set. For example, an input graph data set including a undirected edge that connects a first node N1 and a second node N2 may be received and processed by a graph neural network including a first neuron

NN1 and a second neuron NN2 that are bidirectionally connected to one another, such that message passing can occur from the first node NN1 to the second node NN 2 and, concurrently or consecutively, from the second node NN2 to the first node NN1. An input graph data set including a unidirectional edge that connects a first node N1 to a second node N2 may be received and processed by a graph neural network including a first neuron NN1 (e.g., a neuron in a first layer of a feed-forward graph neural network) that is unidirectionally connected to a second neuron NN2 (e.g., a neuron in a second layer of a feed-forward graph neural network), such that message passing can occur from the first node NN1 to the second node NN 2 but not from the second node NN2 to the first node NN1. An input graph data set including an edge that connects three or more nodes may be received and processed by a graph neural network in which three or more nodes are correspondingly connected.

[3062] Some graph neural networks may be configured to receive and process an input graph data set including a homogeneous set of nodes and/or a homogeneous set of edges. For example, a first neuron of the graph neural network that corresponds to a first node and/or edge of the input graph data set may include a same or similar number of inputs, a same or similar activation function, and/or a same or similar number of outputs as a second neuron of the graph neural network that corresponds to a second node and/or edge of the input graph data set.

[3063] Some graph neural networks may be configured to receive and process an input graph data set including a heterogeneous set of nodes and/or a heterogeneous set of edges. For example, different nodes of an input graph data set may be associated with different labels that respectively indicate different classifications of the nodes, and/or different edges of the input graph data set may be associated with different labels that respectively indicate different classifications of the edges. An architecture of the graph neural network may exhibit variations corresponding to the heterogeneity of the nodes and/or edges. For example, a first neuron of the graph neural network that corresponds to a first node and/or edge of the input graph data set that is associated with a first label or classification may include a different number of inputs, a different activation function, and/or a different number of outputs as a second neuron of the graph neural network that corresponds to a second node and/or edge of the input graph data set that is associated with a second label or classification. As another example, a graph neural network may include a first layer that receives and processes, as input, a first portion of an input data set that includes a first subset of neurons and/or edges that are associated with a first label or classification, and a second layer that receives and processes, as input, a second portion of an input data set that includes a second subset of neurons and/or edges that are associated with a second label or classification. The first layer and the second layer may be processed concurrently or consecutively. The first layer and the second layer may be processed independently (e.g., each layer providing a different portion of an output graph data set). Alternatively, the first layer and the second layer may be processed together (e.g., an output of the first layer may be additionally provided as input to the second layer, and/or an output of the second layer may be additionally provided as input to the first layer).

[3064] Some graph neural networks may include an architecture that is based on one or more node properties of one or more nodes of an input graph data set, one or more edge properties of one or

more edges of the input graph data set, and/or one or more graph properties of the input graph data set. As an example, in some input graph data sets, one or more nodes may include a node property indicating a weight of the node (e.g., an indication of a centrality and/or betweenness of a node among at least a portion of the nodes of the input graph data set). The graph neural network may include a neuron that corresponds to the node, wherein one or more weights of synapses that connect the neuron to other neurons of the graph neural network is based on the weight of the node. As another example, in some input graph data sets, one or more edges may include an edge property indicating a weight of the edge (e.g., an indication of a significance and/or priority of a relationship among two or more nodes of the input graph data set). The graph neural network may include two or more nodes that are connected by a synapse, wherein a weight of the synapse connecting the two or more nodes is based on a weight of an edge of the input graph data set. Examples of node-based graph neural networks include, without limitation, GraphSAGE, PinSAGE, and VR-GCN. Examples of layer-based graph neural networks include, without limitation, FastGCN and LADIES. Examples of subgraph-based graph neural networks include, without limitation, ClusterGCN and GraphSAINT.

[3065] Some graph neural networks may be configured to receive and process fixed input graph data sets, wherein a number and arrangement of nodes and edges of an input data set that is received and processed by the graph neural network does not vary for different instances of processing the input data set. The architecture of such graph neural networks may be configured based on the invariance of the input graph data set. For example, the graph neural network may feature a fixed number and/or arrangement of neurons and/or layers, wherein the fixed architecture of the graph neural network corresponds to the fixed nature of the input graph data set.

[3066] Some graph neural networks may be configured to receive and process dynamic input graph data sets, wherein a number and arrangement of nodes and edges of an input data set that is received and processed by the graph neural network during a first instance of processing can differ from a number and arrangement of nodes and edges of an input data set that is received and processed by the graph neural network during a second instance of processing. As an example, a graph neural network may be configured to perform node and/or edge discovery of an input graph data set and to generate, as output, an output graph data set that includes at least one more node and/or at least one more edge than the input graph data set. Further, the graph neural network may be configured to receive the output graph data set from a first processing as input for a second processing, wherein a number of nodes and/or edges received as input during the second processing is greater than a corresponding number of nodes and/or edges received as input during the first processing. In such cases, an architecture of such graph neural networks may be fixed, but may be configured to receive and process a variety of different input graph data sets (e.g., input graph data sets with a variable number of nodes and/or connections). For example, the graph neural network may include an input layer featuring N input neurons, each corresponding to a node of an input graph data set. Such a graph neural network may be configured to use the fixed architecture to receive and process input graph data sets featuring a variable number of nodes up to, but not exceeding, N. For example, in order to receive and process an input graph data set featuring fewer than N nodes, the graph neural

network may activate only a number of input neurons of the input layer that correspond to the number of nodes in the input graph data set, and to deactivate remaining neurons of the input layer that do not correspond to a node of the input graph data set (e.g., refraining from processing the remaining neurons, and/or processing the neurons but zeroing the weights of the synapses that connect the neurons to other neurons of the graph neural network). As another example, the graph neural network may perform a first processing of a first input graph data set including N nodes, and, accordingly, may deactivate one or more neurons of the input layer. The graph neural network may then perform a second processing of a second input graph data set including more than N nodes (e.g., an output of the first processing may include an output graph data set that includes one or more newly discovered nodes). During the second processing, the graph neural network may activate one or more of the previously deactivated neurons of the input layer in order to receive and process input from the additional nodes of the second input graph data set. For example, the graph neural network may enable or reenable the processing of one or more neurons of the input layer, and/or may reset (e.g., restore and/or initialize) the weights of one or more synapses that connect one or more neurons of the input layer to other neurons of the graph neural network. In some cases, an architecture of such graph neural networks may dynamic, and may change in correspondence with a dynamic nature of the input graph data set. For example, a graph neural network may include an input layer with a variable number of neurons, and may select, adapt, and/or change the number of neurons in the input layer based on a dynamic property of an input graph data set (e.g., a number of nodes and/or edges in the input graph data set). Such graph neural networks may generate new neurons of the input layer (e.g., initializing and/or selecting weights of the synapses of the new neurons, such as copying the weights from the synapses of other neurons of the input layer) based on a larger number of nodes and/or edges of an input graph data set to be received and processed as input. Alternatively or additionally, such graph neural networks may be configured to eliminate and/or merge neurons of the input layer (e.g., initializing and/or selecting weights of the new neurons) based on a smaller number of nodes and/or edges of an input graph data set to be received and processed as input.

[3067] In some graph neural networks, an architecture of the neural network may be selected and/or adapted based on a topology of one or more input graph data sets and/or output graph data sets. For example, a bipartite input graph data set may include two more subgraphs, and a graph neural network may include two or more distinct subsets of neurons that are respectively configured to receive and process data associated with the nodes and/or edges included in one of the subgraphs. As another example, a multigraph input graph data set may include a plurality of edges connecting two or more nodes. For example, a graph representing a social network may include various types of edges that represent various types of relationships (e.g., familial relationships, friendships, and/or professional relationships), and two or more nodes may be connected by a plurality of edges (e.g., a first edge indicating a friendship among the two or more nodes and a second edge indicating a professional relationship among the two or more nodes). An architecture of the graph neural network may correspond to the multigraph nature of the input graph data set. For example, a graph neural network may include two or more distinct subsets of

neurons that are respectively configured to receive and process data associated with a subset of edges of the input graph data set that are of a particular edge type (e.g., a first subset of neurons that is configured to receive and process nodes connected by edges that represent friendships, and a second subset of neurons that is configured to receive and process nodes connected by edges representing professional relationships). As yet another example, an input hypergraph data set may include one or more hyperedges that interconnect three or more nodes. An architecture of a graph neural network that is configured to receive and process the input hypergraph data set may include one or more neurons with synapses that interconnect to two or more other neurons in correspondence with one or more hyperedges of the input hypergraph data set.

[3068] As another example, an architecture of some graph neural networks include one or more layers that perform particular functions on the output of neurons of another layer, such as a pooling layer that performs a pooling operation (e.g., a minimum, a maximum, or an average) of the outputs of one or more neurons, and that generates output that is received by one or more other neurons (e.g., one or more neurons in a following layer of the graph neural network) and/or as an output of the graph neural network. Examples of graph neural networks that include one or more direct pooling layers include, without limitation, SimplePooling, Set2Set, and SortPooling. Examples of graph neural networks that include one or more hierarchical pooling layers include, without limitation, Coarsening, ECC, DiffPool, TopK, gPool, Eigenpooling, and SAGPool.

[3069] As another example, some graph neural networks (e.g., graph convolution networks) include one or more convolutional layers, each of which performs a convolution operation to an output of neurons of a preceding layer of the graph neural network.

[3070] As another example, an architecture of some graph neural networks include memory based on an internal state, wherein the processing of a first input data set causes the graph neural network to generate and/or alter an internal state, and the internal state resulting from the processing of one or more earlier input data sets affects the processing of second and later input data sets. That is, the internal state retains a memory of some aspects of earlier processing that contribute to later processing of the graph neural network. Examples of graph neural networks that include memory features and/or stateful features include graph neural networks featuring one or more gated recurrence units (GRUs) and/or one or more long-short-term-memory (LSTM) cells. In some graph neural networks, these features may be further adapted to accommodate graph processing, such as gated graph neural networks (GGRUs), tree LSTM networks, graph LSTM networks, and/or sentence LSTM networks.

[3071] As another example, an architecture of some graph neural networks includes one or more recurrent and/or reentrant properties. For example, at least a portion of output of the graph neural network during a first processing is included as input to the graph neural network during a second or later processing, and/or at least a portion of an output from a layer is provided as input to the same layer or a preceding layer of the graph neural network. As another example, in some graph neural networks, an output of a neuron is also received as input by the same neuron during the same processing of an input and/or a subsequent processing of an input. The output of the neuron may be evaluated (e.g., weighted, such as decayed) before being provided to the neuron as input.

[3072] As another example, an architecture of some graph neural networks includes two or more subnetworks (e.g., two or more graph neural networks that are configured to process graph data concurrently and/or consecutively). Some graph neural networks include, or are included in, an ensemble of two or more neural networks of the same, similar, or different types (e.g., a graph neural network that outputs data that is processed by a non-graph neural network, Gaussian classifier, random forest, or the like). For example, a random graph forest may include a multitude of graph neural networks, each configured to receive at least a portion of an input graph data set and to generate an output based on a different feature set, different architectures, and/or different forms of processing. The outputs of respective graphs of the random graph forest may be combined in various ways (e.g., a selection of an output based on a minimization and/or maximization of an objective function, or a sum and/or averaging of the outputs) to generate an output of the random graph forest.

[3073] In some cases, an architecture of a graph neural network may be designed by a user. For example, a user may choose one or more hyperparameters of a graph neural network (e.g., a number of layers, a number of neurons in each layer, an activation function used by at least some neurons, and the like) in order to process an input graph data set. In some cases, the selected one or more hyperparameters may be based on domain-specific knowledge, e.g., a specific data type, internal organization or structure, and/or task associated with an input graph data set.

[3074] Alternatively or additionally, in some cases, an architecture of a graph neural network may be selected by an automated process. For example, a hyperparameter search process may determine one or more hyperparameters of a graph neural network based on an analysis of an input graph data set to be received and processed by the graph neural network and/or an analysis of an output graph data set to be generated and provided as output by the graph neural network. The hyperparameter search process may determine various combinations of hyperparameters for variations of the graph neural network (e.g., graph neural networks with different numbers of layers, different numbers of neurons within each layer, graph neural networks including neurons with different activation functions, and/or graph neural networks with different sets of synapses interconnecting the neurons of various layers). The hyperparameter search process may process an input graph data set (e.g., a training input graph data set) using different graph neural networks that correspond to different sets of hyperparameters. The hyperparameter search process may compare the output of the different graph neural networks (e.g., determining a performance measurement for the output of each graph neural network, and comparing the performance measurements of the different graph neural networks) in order to determine and select a graph neural network that generates desirable output (e.g., output that most closely corresponds to a target output associated with the training input graph data set). The hyperparameter search process may discard the other graph neural networks and may use the selected graph neural network to process input graph data sets. In some cases, the hyperparameter search process may iteratively generate and test refined combinations of hyperparameters. For example, after selecting a graph neural network in a first hyperparameter search processing the hyperparameter search process may perform a second hyperparameter search processing by generating additional graph neural networks based on

combinations of hyperparameters that are closer to the hyperparameters of the selected graph neural network, and evaluating the output of the additional graph neural networks. In some cases, the hyperparameter search process may perform a grid search over the set of valid hyperparameter combinations. Iterative refinement of the hyperparameters may enable the hyperparameter search process to determine an architecture of a graph neural network that is well-tuned to a particular task (e.g., an architecture of a graph neural network that demonstrates consistently high performance on input graph data sets within a particular domain of data and/or a particular task). In some cases, a hyperparameter search process may communicate with a user to determine combinations of hyperparameters to evaluate and/or to select for the graph neural network. For example, the hyperparameter search process may present, to a user, a result of a first hyperparameter evaluation (e.g., an output of a graph neural network that was selected through a first hyperparameter search processing). Based on an evaluation of the output by the user, the hyperparameter search process may perform a second or further hyperparameter search processing (e.g., choosing a small refinement of the hyperparameters based on a positive response of the user to the output of a selected graph neural network, and/or choosing a larger refinement of the hyperparameters based on a negative response of the user to the output of the selected graph neural network).

[3075] As another example, some graph neural networks include architectures based on graph convolutional networks (GCNs), wherein a convolutional layer applies a convolution operation to outputs of one or more filters of a previous filter layer of the graph convolutional network. Graph convolutional networks may include spectral convolutional networks that are configured to receive, as input, a spectral representation of an input graph data set, and to apply processing (including one or more convolutional operations) to various spectral components of the spectral representation of the input graph data set. Examples of spectral convolutional networks include, without limitation ChebNet and diversified graph convolutional networks (DGCNs). As another example, some graph convolutional networks include architectures based on spatial convolutional networks (SCNs) that are configured to receive, as input, spatial representations of an input graph data set (e.g., spatial information that represents one or more neighborhoods of nodes and/or edges of the input graph data set), and to apply processing (including one or more convolutional operations) to various spatial components of the spatial representation of the input graph data set. Examples of spatial convolutional networks include, without limitation, spatial convolutional neural networks (SCNNs), spatial and/or spatial-temporal GraphSAGE networks, and some deep convolutional neural networks (DCNNs).

[3076] Graph neural networks can be generated by a variety of machine learning platforms, frameworks, and/or tools, including, without limitation, PyTorch Geometroc, Deep Graph Library, TensorFlow GNN, Graph Nets, Spektral, and Jraph. Frameworks for graph convolutional networks include, without limitation, message passing neural networks (MPNNs), non-local neural networks (NLNNs), mixture model neural networks (MoNet), and Graph Networks (GN).

[3077] Further explanation and/or examples of various architectures of graph neural networks, including the design and implementation of designs and architectures of such graph neural

networks, are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

Training and Performance Evaluation

[3078] Like other types of neural networks, graph neural networks are typically generated with arbitrarily selected parameters (e.g., synaptic weights that are initially set to randomized values). Also, like other types of neural networks, an initialized graph neural network to evaluate input graph data sets through training, in which the parameters of the graph neural network are adjusted to promote desirable processing that produces expected and/or desirable outputs.

[3079] The training of graph neural networks may involve one or more training data sets. For graph neural networks that receive and process input graph data sets, the training data may include one or more training input graph data sets. Alternatively or additionally, for graph neural networks that receive and process input non-graph data, the training data may include one or more sets of training non-graph data.

[3080] The training data for a graph neural network may be based on authentic input data that was previously collected and/or analyzed, or that was collected and analyzed for the purpose of training the graph neural network. For example, in order to process graphs that represent an industrial environment, the training data may include sensor data that was previously and/or is currently received from one or more sensors associated with the industrial environment. Alternatively or additionally, the training data may include partially and/or fully synthetic data. For example, a first portion of training data may include data derived from an analysis of authentic data; authentic data that has been supplemented with synthetic data (e.g., an image of a real-world scene including an inserted artificial object); authentic data that has been modified by a suer (e.g., an image of a real-world scene that has been modified by a user); and/or data generated by one or more algorithms (e.g., other machine learning models and/or simulations of real-world processes). In some cases, the training data set may include both authentic training data and synthetic training data that is based on the authentic training data (e.g., both a real-world image and a modified version of the real-world image that has been adjusted in brightness, contrast, size, resolution, scale, shape, aspect ratio, color depth, or the like).

[3081] The training data for a graph neural network may be limited to a selected data domain. For example, training data for a graph neural network that analyzes social networks may include one or more samples of individuals from within one or more selected social networks. In other cases, the training data for a graph neural network may be generated from a variety of data domains. For example, training data for a graph neural network that analyzes geographic data may include one or more samples of locations of interest and interconnecting pathways from natural outdoor geographic regions (e.g., forests), artificial outdoor geographic regions (e.g., road networks), indoor geographic regions (e.g., caves or shopping malls), historic geographic regions (e.g., maps from ancestral eras and/or civilizations), and/or synthetic geographic regions (e.g., geographic maps from videogames).

[3082] The training data for a graph neural network may be wholly or partially unlabeled. For example, the training data set for an industrial environment may include sensor measurements

collected from the industrial environment, but may not include any data indicating an analysis, classification, metadata, interpolations, extrapolations, interpretation, explanation, and/or user reaction associated with the sensor measurements. Alternatively or additionally, the training data for a graph neural network may be wholly or partially labeled. For example, the training data set for an industrial environment may include sensor measurements collected from the industrial environment, and one or more subsets of sensor measurements may be associated with one or more analyses, classification labels, metadata, interpolations, extrapolations, determinations, interpretations, explanations, and/or user reactions associated with the subset of sensor measurements. Training data may associate labels, metadata, or the like with one or more nodes and/or node properties of a training input graph data set; one or more edges and/or edge properties of a training input graph data set; one or more graph properties of the training input graph data set; and/or one or more portions of non-graph data of a training input data set. In some cases, the labels, data, metadata, or the like associated with at least a portion of a training input data set are selected by one or more users (e.g., a human classification of at least a portion of the training data set). In some cases, the labels, data, metadata, or the like associated with at least a portion of a training input data set are selected by another algorithm (e.g., a simulation or another machine learning model). In some cases, the labels, data, metadata, or the like associated with at least a portion of a training input data set are selected by a cooperation of a human and an algorithm (e.g., a determination by a simulation or another machine learning model that is verified by a reviewing human user).

[3083] Graph neural networks can be trained based on one or more training data sets and one or more learning techniques. As an example, some graph neural networks are trained through an unsupervised learning technique. For example, a training input data set may not include any labels, data, metadata, or the like associated with various portions of the training input data set. The graph neural network may be trained to identify patterns arising within the training input data sets. For example, a training input data set may include data that indicates one or more anomalies (e.g., nodes and/or edges that appear to represent outliers in a data distribution of the nodes and/or edges of the graph) and/or distinctive patterns or structures arising in the data (e.g., cycles arising in a directed and/or undirected graph). The graph neural network may be trained to detect such anomalies, patterns, and/or structure in the training input data sets. The results of unsupervised learning of a graph neural network may be evaluated based on an evaluation of the output of the graph neural network (e.g., a confusion matrix that includes determinations of true positive determinations, true negative determinations, false positive determinations, and/or false negative determinations) and/or performance scores (e.g., an F1 performance score based on ratios of true positives, false positives, true negatives, and false negatives). The weights of various parameters of the graph neural network can be automatically adjusted, corrected, refined, or the like, such that subsequent processing of the same input training data set and/or other input training data sets generates improved evaluations and/or performance scores.

[3084] As another example, some graph neural networks are trained through a supervised learning technique. For example, a training input data set may associate respective portions (e.g., respective

training data samples, such as different training input graph data sets) with one or more labeled outputs that are expected and/or desirable of the trained graph neural network. As an example, a graph neural network may be trained to output a classification of a training input graph data set and/or one or more nodes and/or edges thereof. During a supervised learning process, the training input graph data set may be provided as input to the graph neural network and processed by the graph neural network to generate a predicted classification of a training input graph data set and/or one or more nodes and/or edges thereof. The predicted classifications may be compared with one or more labeled outputs associated with the training input graph data set (e.g., one or more labels associated with an expected and/or desirable classification of the training input graph data set and/or one or more nodes and/or edges thereof). Based on the comparison, the weights of various parameters of the graph neural network can be automatically adjusted, corrected, refined, or the like, such that subsequent processing of the same input training data set and/or other input training data sets generates improved evaluations and/or performance scores (e.g., more accurate predictions of one or more labels associated with an expected and/or desirable classification of the training input graph data set and/or one or more nodes and/or edges thereof).. As another example, a graph neural network may be trained to generate, as output, an output graph data set that is based on a processing of a training input graph data set. During a supervised learning process, the training input graph data set may be provided as input to the graph neural network and processed by the graph neural network to generate an output graph data set. The output graph data set generated by the graph neural network may be compared with one or more expected and/or desirable output graph data sets corresponding to the training input graph data set (e.g., one or more output graph data sets that are expected and/or desired as output when the graph neural network processes the training input graph data set). Based on the comparison, the weights of various parameters of the graph neural network can be automatically adjusted, corrected, refined, or the like, such that subsequent processing of the same input training data set and/or other input training data sets generates improved evaluations and/or performance scores (e.g., more desirable and/or expected output graph data sets).

[3085] As another example, some graph neural networks are trained through a blended training process that includes both supervised and unsupervised learning. For example, a blended training process may evaluate the performance of a graph neural network in training based on both a comparison of predicted outputs of the graph neural network to expected and/or desirable outputs corresponding to an input training data set, and based on one or more automatically determined performance metrics, such as a confusion matrix and/or F1 scores. Some blended training processes may include a round of supervised learning following by a round of unsupervised learning, or may perform rounds of training that include both supervised and unsupervised learning techniques (e.g., optionally with different weights and/or performance thresholds associated with the evaluation of the graph neural network and the updating of the parameters).

[3086] As another example, some graph neural networks are trained through a semi-supervised learning process. For example, a training data set may include a large number of samples, of which only a small number of samples are labeled (e.g., associated with expected and/or desirable

outputs) and a large remainder of the samples are unlabeled (e.g., not associated with expected and/or desirable outputs). The graph neural network may be trained based on the labeled and/or unlabeled training data, and a performance of the graph neural network may be evaluated based on the labels and/or other metrics. In particular, some unlabeled portions of the input training data may be identified as being incorrectly evaluated by the graph neural network (e.g., the graph neural network may generate incorrect outputs such as predictions or classifications, incorrect and/or malformed output graph data sets, or the like). At least a portion of such unlabeled portions of the input training data (e.g., training data samples that appear to be difficult to classify correctly and/or with high confidence) may be submitted to a human reviewer, and the semi-supervised learning process may receive, from the human reviewer, one or more labels that correspond to an expected and/or desirable output of the graph neural network for such portions of the input training data. Training or retraining of the graph neural network may involve the newly labeled portions of the input training data, as well as other portions of the input training data. Semi-supervised learning may enable graph neural networks to be trained based on a smaller degree of human involvement (e.g., a smaller number of labels associated with portions of the input training data set by human reviewers), and may therefore improve a speed, cost, and/or performance of training the graph neural network.

[3087] A training of a graph neural network may occur in one or more epochs. For example, for each epoch, the graph neural network may be provided with input comprising each portion of a training data set, a performance of the graph neural network may be determined based on the output of the graph neural network for each portion of the training data set. Based on the determined performance, and one or more parameters of the graph neural network may be updated. For example, weights of the synapses between neurons of the graph neural network may be adjusted such that a performance of the graph neural network over each portion of the training data set. During the training of a graph neural network, various techniques may be used to evaluate the performance of the graph neural network. As a first example, outputs of the graph neural network (e.g., output graph data sets and/or predictions, such as classifications of the graph, one or more nodes, and/or one or more edges) may be compared with expected and/or desirable outputs. Differences between the outputs and the expected and/or desirable outputs may be used to determine an entropy and/or loss of the output of the graph neural network as compared with corresponding expected and/or desirable outputs. In some variations, the entropy or loss of the graph neural network determined during or after a current epoch may be compared with an entropy or loss of the graph neural network determined during or after a previous epoch to determine a differential and/or marginal entropy or loss. A negative differential and/or marginal entropy or loss may indicate that the training of the graph neural network is productive (e.g., the performance of the graph neural network improved in the current epoch as compared with a previous epoch). A zero or positive differential and/or marginal entropy or loss may indicate that the training of the graph neural network is unproductive (e.g., the performance of the graph neural network did not improve, or diminished, in the current epoch as compared with a previous epoch). Training of the graph neural network may therefore continue as long as the differential and/or marginal entropy or

loss remains negative and, optionally, exceeds a threshold magnitude that indicates significant training progress.

[3088] As another example, outputs of the graph neural network (e.g., output graph data sets and/or predictions, such as classifications of the graph, one or more nodes, and/or one or more edges) may be classified as one of a true positive, a false positive, a true negative, or a false negative. The performance of the graph neural network may be evaluated as a confusion matrix, e.g., based on a calculation of the performance over the incidence of true positive, false positive, true negative and false negative outputs. In some cases, the calculation may be weighted based on a risk matrix that applies different weights to each classification of the output. For example, in a graph neural network that generates classifications of graphs that correspond to diagnoses of medical conditions, it may be determined false negatives (e.g., missed diagnoses) are very harmful or costly, while false positives (e.g., misdiagnoses that can be corrected by further evaluation) may be determined to be comparatively harmless. Accordingly, the performance of the graph neural network may be determined based on a weighted calculation over the confusion matrix that more severely penalizes the performance based on false negatives than false positives.

[3089] As another example, the training of a graph neural network may involve an improvement of an objective function that serves as a basis for measuring the performance of the graph neural network. For example, the objective function may include (without limitation) a loss minimization, an entropy minimization, a precision maximization, a recall maximization, an error minimization, or a consistency maximization. The objective function may include a comparison of the performance of the graph neural network over various distributions of the input data set (e.g., a minimax optimization, such as minimizing a maximum loss over any portion of the input data set, or a maximin optimization, such as maximizing a minimum loss over any portion of the input data set). In some training scenarios that involve reinforcement learning, the output of a graph neural network may include and/or may be interpreted as a policy, e.g., a set of responses of an agent based on respective conditions. The performance of the graph neural network may be based on various objective functions that evaluate various properties of the generated and/or interpreted policy. For example, in a q-learning reinforcement learning process, the objective function applied to the policy may include a maximization of an action value of each behavior that may be performed in response to various conditions.

[3090] As another example, the training of graph neural networks may occur concurrently with the hyperparameter search and/or selection. For example, a hyperparameter search process may initially identify a first set of combinations of hyperparameters of graph neural networks to be evaluated using a training data set. Based on each such combination of hyperparameters, a graph neural network may be generated and at least partially trained to determine its performance. Based on the evaluation of the outputs of the graph neural networks corresponding to respective combinations of hyperparameters, the hyperparameter search process may identify a candidate graph neural network with a highest performance. The hyperparameter search process may then generate a second set of combinations of hyperparameters based on the hyperparameters of the candidate graph neural network, and may further (at least partially) train and evaluate the

performance of additional graph neural networks based on the second set of combinations of hyperparameters. A comparison of the performance of the additional graph neural networks may cause the hyperparameter search process to retain the candidate graph neural network or to choose a new candidate graph neural network from among the additional graph neural networks. The hyperparameter search process may continue until additional improvements in the performance of candidate graph neural networks are not achievable and/or are below a threshold performance improvement. In this selection process, a variety of performance metrics may be used. As previously discussed, the performance metrics may include an evaluation of the outputs of the graph neural networks (e.g., a loss or entropy, a differential or marginal loss or entropy, a confusion matrix, an F1 score, or the like). Alternatively or additionally, the performance metrics may include other features of the output, such as a consistency of the output of the graph neural network over the distribution of data in the training data set and/or a bias in the performance the output of the graph neural network for selected data distributions of the training data set, and/or a smoothness or oversmoothness of the graph nodes represented in the graph neural network. Alternatively or additionally, the performance metrics may include one or more measurements of computational resource expenditures to perform training and/or inference of input data sets with the graph neural network (e.g., CPU and/or GPU utilization, memory usage, training time and/or complexity, processing latency between receiving input and generating output, or the like). Aggregate performance measurements may be based on a variety of such considerations, and may enable a human designer and/or a hyperparameter search process to perform a selection of a graph neural network based on various performance tradeoffs (e.g., a preference for a first graph neural network that produces high-accuracy, high-consistency, and/or high-confidence results but that requires a large amount of computational resources, time, and/or cost, vs. a preference for a second graph neural network that produces reasonable-accuracy, reasonable-consistency, and/or reasonable-confidence results using a smaller amount of computational resources, time, and/or cost). For example, a measurement of computational resource utilization by a particular graph neural network may correspond to a numeric penalty in various measurement of the performance of the graph neural network (e.g., a loss, entropy, and/or objective function output).

[3091] In various forms of graph neural network training based on these and other learning techniques, various training methods can be used to update the parameters of a graph neural network in training and/or to evaluate the performance of a graph neural network in training. For example, optimizers that may be used during the training of graph neural networks may include (without limitation) linear regression; root mean squared propagation (RMSprop); stochastic gradient descent; adaptive stochastic gradient descent (Adagrad); adaptive stochastic gradient descent with adaptive learning (Adadelta); adaptive moment estimation (Adam); Nesterov accelerated adaptive moment estimation (Nadam); Nesterov accelerated gradient and momentum (NAG); Monte Carlo simulations involving various variance reduction techniques, such as control variates; or the like, including variations and/or combinations thereof. Training techniques for particular types of graph neural networks may include optimizers that are specialized for such

particular types of graph neural networks (e.g., graph convolutional networks may be trained using a FastGCN optimizer and/or receptive field control (RFC) optimizers).

[3092] As further examples, graph neural network training may include a variety of techniques that are also applicable to non-graph machine learning models, including non-graph neural networks. As a first such example, training may occur in batches and/or mini-batches of the training data set, wherein the graph neural network evaluates a batch (e.g., plurality of input data sets) of an input training data set, and the parameters of the graph neural network are updated based on an aggregation of the evaluation of the outputs of the graph neural network for the batch of input data sets. In various training techniques, batches may be selected at random from the training input data set or may be selected in an organized manner, e.g., as various subsets that are representative of one or more data distributions of the training input data set. For example, if the graph neural network in training exhibits good performance over some data distributions of the training input data and poor performance over other data distributions of the training input data, the continued training of the graph neural network may focus on, prioritize, and/or overweight the training based on batches of training input data that reflect the data distributions associated with poor performance. In various training techniques, a batch size of batches of training input data sets may be fixed, or the batch size may vary based on the progress of the training of the graph neural network.

[3093] As another example, in various training techniques for graph neural networks, an entire set of training input data may be partitioned into a training data set that is used only to train the graph neural network and update its parameters; a validation data set that is used only to evaluate a prospective and/or in-training graph neural network; and/or a test data set that is used to only evaluate a final performance of the fully trained graph neural network. The partitioning of the training input data may be based on one or more ratios (e.g., a 90/5/5 partitioning of the training input data into a training data set, a validation data set, and a test data set, or a 98/1/1 partitioning of the training input data into a training data set, a validation data set, and a test data set). For example, during an epoch, the performance of the graph neural network may be evaluated based on various portions of the training data set, and the parameters of the graph neural network may be adjusted based on the determined performance. However, continued training and updating of the graph neural network based on the training data set may result in overfitting, e.g., "memoization" of correct outputs that correspond to various portions of the training data set. Due to such overfitting, the performance of the graph neural network in evaluating previously evaluated input data sets may improve, but performance of the graph neural network on previously unevaluated input data sets may decline. Instead, at the conclusion of an epoch, the performance of the graph neural network may instead be evaluated based on various portions of the validation data set, which is not otherwise used to update the parameters of the graph neural network. Evaluation of the performance of the graph neural network on previously unseen data can indicate that the performance of the graph neural network is genuinely improving (e.g., based on learned principles of data evaluation that apply consistently to both previously seen and previously unseen input data sets), resulting in a continuation of training. Alternatively, Evaluation of the performance of the

graph neural network on previously unseen data can indicate that the performance of the graph neural network is resulting in overfitting to the training data set (e.g., based on "memoization" of correct outputs for previously seen input data sets that do not inform the correct evaluation of previously unseen input data sets), resulting in a conclusion of training. Such conclusion may be referred to as "early stopping" of training to reduce overfitting of the graph neural network to the training data set and to preserve the performance of the graph neural network on previously unsee input data sets.

[3094] As another example, various training techniques for graph neural networks may include one or more regularization techniques, in which the inputs to the graph neural network and/or the processing of the input are adjusted to reduce overfitting. As a first example, the training of a graph neural network may include a dropout regularization technique, in which some neurons of the graph neural network are disabled for some instances of processing input data sets. In various regularization techniques, neurons to be disabled are selected randomly (e.g., 5% of the neurons during each epoch) and/or can be selected in a sequence (e.g., a round-robin selection of deactivated neurons). The selected neurons may be disabled by refraining from processing the inputs of the neurons and setting the outputs of the selected neurons to zero, and/or by processing the selected neurons but temporarily setting the weights of the synapses of the neurons to zero. As a second example, the training of a graph neural network may include a dropnode and/or dropedge regularization technique, in which portions of an input graph data set that include some nodes and/or some edges of the input graph data set are disabled. In various regularization techniques, nodes and/or edges to be disabled for an instance of processing are selected randomly (e.g., 5% of the nodes and/or edges during each epoch) and/or can be selected in a sequence (e.g., a round-robin selection of deactivated nodes and/or edges). The selected nodes and/or edges may be disabled by refraining from processing portions of the input data set that correspond to the selected nodes and/or edges, and/or by deactivating neurons of an input layer of the graph neural network that are configured to receive input data from the selected nodes and/or edges. As a third example, the performance of a graph neural network may be subjected to various forms of regularization, including L1 ("lasso") regularization and/or L2 ("ridge") regularization. These and other forms of regularization may be used, alone or in combination, to reduce overfitting of a graph neural network to an input training data set. For example, regularization may reduce an overweighting of a subset of nodes, edges, and/or neurons in the processing of various input data sets (e.g., by reducing and/or penalizing neurons having synaptic weights with magnitudes that are disproportionately large compared to the synaptic weights of other neurons of the graph neural network).

[3095] As another example, various training techniques for graph neural networks may combine a graph neural network with one or more other machine learning models, including one or more other graph neural networks and/or one or more non-graph neural networks. For example, a bootstrap aggregation ("bagging") training technique involves a determination of a decision tree as an ensemble of machine learning models based on different bootstrap samples of the training input data set. Each machine learning model, including one or more graph neural networks, may be

trained based on a random subsample of the training input data set. For a particular input data set, many of the trained machine learning models of the ensemble, including one or more graph neural networks, may present poor or only adequate performance. However, one or a few of the trained machine learning models may generate high-performance output for the particular input data set and others like it (e.g., for input data sets that share one or more properties, such as a select graph property, a select node property, and/or a select edge property). Thus, for any particular input data set, an evaluation of the specific properties of the particular input data set may enable a selection among the available models of the ensemble that may be used to evaluate the particular input data set. That is, a machine learning model (e.g., a graph neural network) that is generally a poorly performing model on most input data sets may exhibit good performance over a small neighborhood of input data sets that includes the particular data set, and may therefore be selected to evaluate the particular data set. Alternatively or additionally, the bootstrap aggregation may involve an evaluation of an input data set by a plurality of machine learning models (optionally including one or more graph neural networks of the ensemble) and a combination of the outputs of the selected machine learning models. In such scenarios, it is possible the individual outputs of the individual machine learning models exhibit poor performance (e.g., incorrect and/or low-confidence classifications of an input data set), but a determination of a consensus over the outputs of the multiple machine learning models may exhibit high performance (e.g., accurate and/or high-confidence classifications of the input data set).

[3096] As another example, various training techniques for graph neural networks may include a boosting ensemble technique, in which an output of a first trained machine learning model (e.g., a first graph neural network) is evaluated by a second trained machine learning model (e.g., a second graph neural network) to predict an accuracy and/or confidence of the prediction of the first trained machine learning model. For example, a first trained graph neural network may be evaluated to determine that it generates accurate and/or high-confidence output for a first group of input data sets (e.g., input graph data sets that include a first graph property, a first node property, and/or a first edge property), but inaccurate and/or low-confidence output for a second group of input data sets (e.g., input graph data sets that include a second graph property, a second node property, and/or a second edge property). A particular input data set may initially be processed by the first trained graph neural network to determine a first output (e.g., an output graph neural network or a prediction, such as a classification). A second trained graph neural network may evaluate the input data set and/or the output of the first graph neural network to predict an accuracy and/or confidence of the first graph neural network over input data sets that resemble the particular input data set. If the second trained graph neural network predicts that the output of the first graph neural network is likely to be of high accuracy and/or confidence, then the second trained graph neural network may provide the output of the first trained graph neural network as its output. However, if the second trained graph neural network predicts that the output of the first graph neural network is likely to be of low accuracy and/or confidence, then the second trained graph neural network may adjust, correct, and/or discard the output of the first trained graph neural network, or preferentially select an output of a different machine learning model (e.g., a third trained graph neural network)

to be provided as output instead of the output of the first trained graph neural network. In such scenarios, it is possible that the individual outputs of the individual machine learning models exhibit poor performance (e.g., incorrect and/or low-confidence classifications of an input data set), but the review and validation of the output of some machine learning models by other machine learning models may enable a determination of a consensus over the outputs of the multiple machine learning models that exhibits high performance (e.g., accurate and/or high-confidence classifications of the input data set).

[3097] As another example, following conclusion of training a graph neural network, the graph neural network may be deployed for use (e.g., transferred to one or more devices, deployed into a production environment, and/or connected to a source of production input data). The performance of the graph neural network over input data sets may continue to be evaluated and monitored to verify that the graph neural network continues to perform well over various inputs. In some cases, the performance of the graph neural network may change between training and deployment. For example, a distribution of production input data processed by the graph neural network may differ from the distribution of training input data that was used to train the graph neural network. Alternatively or additionally, a distribution of production input data may change over time, e.g., between a time of deploying the graph neural network and a later time after such deployment. Such instances of changes in the performance of a fully trained and deployed graph neural network may be referred to as "drift." In some such cases, "drift" may be reduced or eliminated by retraining or continuing training of the graph neural network, e.g., using additional training input data that corresponds to an actual or current distribution of the production input data. Alternatively or additionally, "drift" may be reduced or eliminated by training a substitute graph neural network to replace the initially deployed graph neural network. For example, the substitute graph neural network may include a different set of hyperparameters than the initially deployed graph neural network (e.g., additional layers and/or neurons to provide greater learning capacity; additional regularization techniques to reduce overfitting to the training data set; and/or the inclusion of specialized layers, such as pooling, filtering, memory, and/or attention layers). As another example, the initially deployed graph neural network may be added to an ensemble of other machine learning models, optionally including other graph neural networks, to generate improved outputs (e.g., higher-accuracy predictions) based on a consensus determined over the outputs of a number of machine learning models.

[3098] As another example, the training and/or use of graph neural networks may be susceptible to various forms of adversarial attack. For example, in an adversarial attack scenario, a particularly designed and/or selected input to a graph neural network (an "adversarial input," such as an unusual, malformed, and/or anomalous) may cause the graph neural network to generate output that is incorrect, inconsistent with other outputs, and/or surprising. As an example, in a form of graph modification adversarial attack that may be referred to as node injection poisoning adversarial attack (NIPA), one or more nodes of an input graph data set are selected and/or altered to shift an output of the graph neural network based on the adversarial input (e.g., altering a classification and/or prediction of the input graph data set, or altering an output graph data set

based on the adversarial input graph data set). As another example, in a form of graph modification adversarial attack that may be referred to as an edge perturbing adversarial attack (NIPA), one or more edges of an input graph data set are selected and/or altered to shift an output of the graph neural network based on the adversarial input (e.g., altering a classification and/or prediction of the input graph data set, or altering an output graph data set based on the adversarial input graph data set). As another example, in a training data injection attack, one or more portions of training input data on which a graph neural network is trained are designed and/or altered to alter the training of the graph neural network (e.g., a mislabeling of a particular training data input that causes the graph neural network to misclassify other inputs that correspond to the mislabeled training data input, and/or an injection of data samples into a training data set that alter a data distribution of the training data set upon which the graph neural network is trained). As another example, in a membership inference adversarial attack, properties and/or outputs of a graph data set are evaluated to identify properties of one or more training data inputs on which the graph data set was trained (e.g., an influential property of an input data set that causes the graph data set to select a particular classification for the any input data sets that include the property). As another example, in a property inference adversarial attack, properties and/or outputs of a graph data set are evaluated to identify general properties of training data inputs on which the graph data set was trained (e.g., a distribution of data included in the training data set, which may indicate particular distributions of input data over which the graph neural network was not trained, or over which the graph neural network was incompletely and/or incorrectly trained). As another example, in a model inversion adversarial attack, outputs of a graph neural network are examined to identify properties of corresponding input data sets that cause the graph neural network to generate such outputs.

[3099] Based on these and other forms of adversarial attack, the training and/or evaluation of a graph neural network may be adjusted to protect the graph neural network from such adversarial attack. For example, before an input to a graph neural network is processed, the input may be evaluated and/or classified (e.g., by another machine learning model, including another graph neural network) in order to determine whether the input is adversarial. If so, the graph neural network may refrain from processing the adversarial input, may process the adversarial input in more limited conditions (e.g., processing only a portion of the adversarial input, and/or replacing a malformed or anomalous portion of the adversarial input with a corresponding non-malformed and/or non-anomalous portion). As another example, during processing of an input data set, the internal behavior of the graph neural network may be evaluated and/or classified (e.g., by another machine learning model, including another graph neural network) to determine whether the behavior indicates a processing of adversarial input (e.g., unusual neuron activations, unusual outputs of one or more neurons, and/or updates of internal states of memory units). If so, the processing of the adversarial input may be halted and/or an internal state of the graph neural network may be restored to a time before the adversarial input was processed. As another example, before output of a graph neural network is provided in response to an input data set, the output may be examined and/or classified (e.g., by another machine learning model, including another graph neural network) to determine whether it is incorrect, inconsistent with other inputs, and/or

surprising. If so, the output of the graph neural network may be discarded and/or altered before being provided in response to the input data set. Further explanation and/or examples of various techniques for training and performance evaluation of graph neural networks are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

Applications

[3100] Graph neural networks can be applied to input data sets (including input graph data sets and/or input non-graph data sets) in various applications, and can be configured and/or trained to generate outputs (including output graph data sets and/or output predictions, such as classifications) that are relevant to various tasks within such applications.

[3101] For example, in the field of social networking, a graph data set may represent at least a portion of a social network, including nodes that represent people and that are connected by edges that represent relationships among two or more people. The graph data set representing a social network may be provided, as input, to a graph neural network that is configured to receive and process the input graph data set. The graph neural network may generate, as output, an output graph data set. For example, the output graph data set may include one or more new nodes that correspond to one or more newly discovered people within the social network, and/or one or more new edges that correspond to one or more newly discovered relationships that connect two or more people of the social network. The output graph data set may include one or more subgraphs and/or clusters that represent highly interconnected people of the social network, e.g., a social circle. The output graph data set may include a prediction of a recommendation of a relationship among two or more nodes corresponding to two or more people of the social network who share common personal traits, interests, and/or connections to other people. The output graph data set may include a prediction of a classification of a node corresponding to a person of the social network, e.g., a prediction of a personal interest of the person or a demographic trait of the person. The output graph data set may include a prediction of a classification of an edge that connects nodes representing two or more people of the social network, e.g., a prediction of a criminal association among two or more people of the social network. The output graph data set may include a determination of a relationship within the social network based on an attention model, e.g., an identification of a first node corresponding to a first person of the social network that appears to be influential to a second person of the social network represented by a second node of the graph. The output graph data set may include a prediction of a graph property of the graph, e.g., a classification of the social network as one or more types (e.g., a genealogy or familial social network, a friendship social network, and/or a professional relationship social network).

[3102] As another example, in the field of pharmaceuticals, a graph data set may represent at least a portion of a molecule (e.g., a protein or a DNA sequence), including nodes that represent atoms of the molecule and that are connected by edges that represent bonds and/or spatial relationships among two or more atoms. The graph data set representing a molecule may be provided, as input, to a graph neural network that is configured to receive and process the input graph data set. The graph neural network may generate, as output, an output graph data set. For example, the output

graph data set may include one or more new nodes that correspond to one or more newly discovered atoms that may be added to the molecule, and/or one or more new edges that correspond to one or more newly discovered atoms of the molecule. The output graph data set may include one or more subgraphs and/or clusters that represent highly interconnected subregions of the molecule, such as carbon atoms that form a benzene ring or a binding site for a protein. The output graph data set may include a prediction of a classification of one or more nodes corresponding to one or more atoms of the molecule, e.g., a prediction that a subset of atoms of the molecule include a binding site for an enzyme that may active and/or deactivate a protein. The output graph data set may include a prediction of a classification of an edge that connects nodes representing atoms of the molecule, e.g., a prediction of a chemically reactive bond that can be altered to alter a property of the molecule. The output graph data set may include a prediction of a graph property of the graph, e.g., a prediction of a shape or organization of the molecule, a classification of the molecule as an enzyme, and/or a prediction of a potential side-effect of a drug due to an undesirable interaction with another drug.

[3103] As another example, in the field of software, a graph data set may represent at least a portion of a marketplace, including nodes that represent products and that are connected by edges that represent relationships between products. The graph data set representing a marketplace may be provided, as input, to a graph neural network that is configured to receive and process the input graph data set. The graph neural network may generate, as output, an output graph data set. For example, the output graph data set may include one or more new nodes that correspond to one or more newly discovered product, and/or one or more new edges that correspond to one or more newly discovered products. The output graph data set may include one or more subgraphs and/or clusters that represent highly interconnected products (e.g., two or more products that are often purchased and/or used together, or that compete in a particular market sector). The output graph data set may include a prediction of a recommendation of a relationship among two or more nodes corresponding to two or more products. The output graph data set may include a prediction of a classification of a node corresponding to a product, e.g., a prediction of an appeal, value, and/or demand of a product in a particular market segment, such as a particular subset of users. The output graph data set may include a prediction of a classification of an edge that connects nodes representing products, e.g., a prediction of a functional relationship between two or more products. The output graph data set may include a prediction of a graph property of the graph, e.g., a classification of the marketplace as increasing and/or decreasing in terms of supply, demand, size, prognosis, and/or public interest.

[3104] As another example, in the field of civil engineering, a graph data set may represent at least a portion of a geographic region, including nodes that represent locations of interest and that are connected by edges that represent roads. The graph data set representing a geographic region may be provided, as input, to a graph neural network that is configured to receive and process the input graph data set. The graph neural network may generate, as output, an output graph data set. For example, the output graph data set may include one or more new nodes that correspond to one or more newly discovered locations of interest, and/or one or more new edges that correspond to one

or more newly discovered locations of interest. The output graph data set may include one or more subgraphs and/or clusters that represent highly interconnected locations of interest. The output graph data set may include a prediction of a recommendation of a relationship among two or more nodes corresponding to two or more locations of interest. The output graph data set may include a prediction of a classification of a node corresponding to location of interest, e.g., a prediction of a current or future volume of visitors to a location of interest and/or a volume of traffic at or through the location of interest. The output graph data set may include a prediction of a classification of an edge that connects nodes representing locations of interest, e.g., a prediction of a volume of traffic on a road that connects two or more locations of interest. The output graph data set may include a prediction of a graph property of the graph, e.g., a classification of a sufficiency of a road network of the geographic region to support a current or future volume of traffic.

[3105] As another example, in the field of industrial systems, a graph data set may represent at least a portion of an industrial plant, including nodes that represent machines of the industrial plant and that are connected by edges that represent functional relationships among the machines. The graph data set representing the industrial plant may be provided, as input, to a graph neural network that is configured to receive and process the input graph data set. The graph neural network may generate, as output, an output graph data set. For example, the output graph data set may include one or more new nodes that correspond to one or more newly discovered machines, and/or one or more new edges that correspond to one or more newly discovered machines. The output graph data set may include one or more subgraphs and/or clusters that represent highly interconnected machines. The output graph data set may include a prediction of a recommendation of a relationship among two or more nodes corresponding to two or more machines. The output graph data set may include a prediction of a classification of a node corresponding to a machine, e.g., a prediction of a current or future maintenance state of a machine. The output graph data set may include a prediction of a classification of an edge that connects nodes representing machines, e.g., a prediction of a functional relationship between a first machine and a second machine that may significantly impact an efficiency, output, cost, or the like of the industrial plant. The output graph data set may include a prediction of a graph property of the graph, e.g., a classification of the industrial plant as belonging to a particular industry, such as raw material processing, semiconductor fabrication, tool manufacturing, vehicle manufacturing, textile manufacturing, and/or pharmaceuticals manufacturing. The output graph data set may include a prediction of a future and/or optimized state of the industrial plant, e.g., a reorganization of the machines of the industrial plant to optimize machine placement and/or floor planning.

[3106] As another example, in the field of cybersecurity, a graph data set may represent at least a portion of a device network, including nodes that represent devices and that are connected by edges that represent communication and/or interactions among two or more devices. The graph data set representing the device network may be provided, as input, to a graph neural network that is configured to receive and process the input graph data set. The graph neural network may generate, as output, an output graph data set. For example, the output graph data set may include one or more new nodes that correspond to one or more newly discovered devices, and/or one or more new edges

that correspond to one or more newly discovered devices. The output graph data set may include one or more subgraphs and/or clusters that represent highly interconnected devices. The output graph data set may include a prediction of a recommendation of a relationship among two or more nodes corresponding to two or more devices. The output graph data set may include a prediction of a classification of a node corresponding to a device, e.g., a prediction of a security status of the device as being safe, vulnerable, or corrupted. The output graph data set may include a prediction of an activity occurring among the nodes of the graph data set, e.g., an occurrence of an intrusion or an attack based on anomalous activities represented by the edges of the graph data set. The output graph data set may include a prediction of a classification of an edge that connects nodes representing devices, e.g., a prediction that a particular interaction between two or more devices is associated with a security vulnerability or attack. The output graph data set may include a prediction of a graph property of the graph, e.g., a classification of the set of devices as safe from security flaws or vulnerable to one or more attack mechanisms, such as denial-of-service (DoS) attacks, distributed-denial-of-service (DDoS) attacks, social engineering attacks such as phishing, eavesdropping attacks such as man-in-the-middle attacks, or the like. The output graph data set may include a prediction of a theoretical state of the graph data set, e.g., a security state of the device network in response to a particular type of attack, and/or a security state of the device network based on the inclusion of additional devices in the future. The output graph data set may include a recommendation to modify the graph neural network based on one or more security considerations, e.g., a recommendation to reorganize the device network to reduce susceptibilities to one or more security risks. The output graph data set may include a technique to defend the graph neural network from various types of adversarial attack, e.g., training-time attacks that affect the manner in which the graph neural network learns to evaluate and/or classify the graph data set, one or more nodes, and/or one or more edges. For example, the message passing operations of the graph neural network may be modified to reduce a susceptibility of the graph neural network to adversarial perturbation during training, while preserving the learning capabilities of the graph neural network.

[3107] Examples of additional applications of various graph neural networks to various graph data sets include, without limitation: graph mining applications (e.g., graph matching and/or clustering); physics (e.g., physical systems modeling and/or evolution over time); chemistry (e.g., molecular fingerprints and/or chemical reaction predictions); biology (e.g., protein interface predictions, side effects predictions, and/or disease classification); knowledge graphs (e.g., knowledge graph completion and/or knowledge graph alignment); generation (e.g., output graph data set generation that corresponds to an expression, an image, a video, a music sample, or a scene graph); combinatorial optimization; traffic networks (e.g., traffic state prediction); recommendation systems (e.g., user-item interaction predictions and/or social recommendations); economic networks (e.g., stock markets); software and information technology (e.g., software defined networks, AMR graph-to-text tasks, and program verification); text processing (e.g., text classification, sequence labeling, machine translation, relation extraction, event extraction, fact verification, question answering, and/or relational reasoning); and image processing (e.g., social

relationship understanding, image classification, visual question answering, object detection, interaction detection, region classification, and/or semantic segmentation). Further examples of applications for processing various graph data sets by various graph neural networks are presented elsewhere in this disclosure and/or will be known to or appreciated by persons of ordinary skill in the art.

[3108] In example embodiments, graph neural networks may be used in connection with a set of digital twins. For example, graph neural nets may be used to model the relationships between different digital twins, model the relationships between different components of a digital twin, predict the behavior of a digital twin under different conditions, and/or optimize the performance of digital twins.

Transformer Models in Artificial Intelligence Platforms

Attention

[3109] In example embodiments, an artificial intelligence system, machine learning model, or the like, of any of the types disclosed herein, may comprise, integrate, link to, or include an attention feature. Attention may be generally described as a determination, among a set of inputs, of the relatedness of each input to the other inputs in the set of inputs. In "self-attention," the input includes a sequence of elements, and attention is determined between each pair of elements in the sequence. As a first example, the set of inputs includes a sequence of words in a language, and attention is applied to determine, for each word in the sequence, the relatedness of the word to each other word in the sequence. As a second example, an input includes an image comprising a set of pixels, and attention is applied to determine, for each group of pixels in the image, the relatedness of the group of pixels to each other group of pixels in the image. Attention can also be applied between sets of input, wherein attention is determined between each element of a first set of input and each element of a second set of input. For example, the set of inputs can include a first sequence of words in a first language and a second sequence of words in a second language, and attention can be determined to indicate how each word in the first sequence is related to each word in the second sequence.

[3110] FIG. 269 presents an example of a determination of attention by a machine learning model. In the example of FIG. 269, an input sequence includes a set of tokens, each representing a word ("The," "Furry," "Dog," "Chased," "The," "Cat"). Each token includes an indicator of a position of the token in the sequence. In various embodiments, the tokens of the input sequence may include complete words, portions of words (e.g., a first token indicating a word root and a second token indicating a modifier of the word root), punctuation, or the like. Some tokens may indicate metadata, such as a start-of-sequence token, an end-of-sequence token, or a null token indicating a padding of the sequence or a mask that hides a token of the sequence.

[3111] The input sequence is processed by a position encoder that determines, for each token, an encoding of the position. In example embodiments, the position encoding may include an ordinal numerical value that indices the ordinal position of each token in the sequence, such as an index beginning at zero or one. In example embodiments, the position encoding may include a relative numerical value that indicates a position of each token in the sequence relative to a fixed position,

such as a current word (encoded position 0), an immediately preceding word (encoded position -1), or an immediately following word (encoded position 1). In example embodiments, the position encoding may include non-integer values and/or multiple values, such as a first index indicating a sine calculation (with a given frequency) of the position of each token and a second index indicating a cosine calculation (with a same or different frequency) of the position of each token.

[3112] The input sequence is also processed by an embedding model. The embedding model determines, for each token in the input sequence, a mapping of the token into a latent space representation of the input (e.g., a latent space representation of a language). The latent space may position each token along a plurality of n dimensions, wherein each dimension represents a distinct type of relationship among the elements of the language. The embedding model clusters the tokens such that related tokens are positioned closer to each other within the latent space. For example, along one dimension of the latent space, the words "Cat" and "Dog" may be positioned close together as being words that describe animals, while also being positioned apart from words that do not describe animals, such as "Baseball" and "School." Along another dimension of the latent space, the words "Dog" and "Furry" may be positioned close together as words that commonly occur in the context of dogs, while also being positioned apart from words that do not describe dogs, including "Cat." For each token of the input sequence, the embedding model generates one or more values that indicate the position of the token within the latent space. In example embodiments, the values are encoded as a vector, and the proximity of two tokens within the latent space may be determined based on vector proximity calculations, such as cosine similarity.

[3113] Based on the positions encoded by the position encoder and the embeddings determined by the embedding model, a model input can be generated for the input sequence. As shown in FIG. 269, the model input includes a query, a set of keys, and a set of values. As an example, the query may include an indicator of a particular token in the input sequence, such as the sixth token ("Cat"). The keys may include the position encodings of respective tokens of the input sequence, as determined by the position encoder, and a corresponding embedding of the respective token as determined by the embedding model. The values may indicate additional data features of the tokens. As an example, the values may indicate, for each token of the input sequence, a determined sentiment (e.g., a ranking between -1, indicating very negative words, and +1, indicating very positive words). In example embodiments, no additional data features are available, and the values are identical to the keys.

[3114] The model input is received and processed by an attention layer. In FIG. 269, the attention layer first includes a set of fully-connected layers: a first fully-connected layer processes the query of the model input; a second fully-connected layer processes the keys of the model input; and a third fully-connected layer processes the values of the model input. Each fully-connected layer includes a bias and a set of weights that adjust the values of the query, key, or value, respectively. The bias and weights of each fully-connected layer are model parameters that are initialized (e.g., to random values) and then incrementally adjusted during training.

[3115] Optionally, in example embodiments, the outputs of the fully-connected layers are further processed by a masking layer. The masking layer removes one or more values from the model

input adjusted by the fully-connected layers. As a first example, the masking layer can reduce to zero the values of the key and/or value at a given position, such as a token at a current position to be predicted, or a token at a position following the current position that is to be hidden from the model. As a second example, the masking layer can reduce to zero the values of particular keys and/or values, such as padding values that are provided to adapt the size of the model input to a size of input that the attention layer is configured to receive and process. The masking layer can produce output for certain tokens (e.g., reduced to zero) for the indicated tokens (e.g., the current token, future tokens, and/or padding tokens) and that is the same as the input for the remaining tokens.

[3116] Optionally, in example embodiments, the outputs of the masking layer are further processed by a multi-head reshaping layer. The multi-head reshaping layer can reshape an input vector comprising the weighted and/or masked model input such that subsets of the input can be processed in parallel by different attention heads. As an example, an attention layer may include two attention heads, and the input can be reshaped such that each attention head is applied to only half of the inputs. The multi-head attention model can enable attention determinations over different subsets of the input (e.g., a first attention head can determine the relatedness of a first token to a first subset of tokens of the input sequence, and a second attention head can determine the relatedness of the same first token to a second subset of tokens of the input sequence). Alternatively or additionally, the multi-head attention model can enable different types of attention determinations among the tokens of the input sequence (e.g., a first attention head can determine a first type of relatedness of a first token to a subset of tokens of the input sequence, and a second attention head can determine a second type of relatedness of the same first token to the same or different subset of tokens of the input sequence). The multi-head attention model may enable parallel processing of the input sequence (e.g., the input for each attention head can be processed by a different processing core).

[3117] The attention layer includes an attention calculation that determines, based on the model input, the attention of a token of the input sequence with respect to other tokens of the input sequence. In example embodiments, the attention calculation includes an additive attention ("Bahdanau Attention") calculation, in which attention is determined as a sum of weighted calculations of the distances of the tokens along each dimension of the latent space. In example embodiments, the attention calculation includes a dot product determination, as a comparison of the distances between the vectors of the tokens within the latent space. In example embodiments, the attention calculation is performed over the query, keys, and values of the model input, optionally after processing with a masking layer. In example embodiments, the attention calculation is performed for each of a plurality of attention heads, each of which processes a particular subset of the tokens of the input sequence.

[3118] In embodiments that include multi-head reshaping, the output of the attention calculation is further processed by a merge operation that merges the attention calculations for the respective attention heads. In example embodiments, the merge operation includes a concatenation and/or interleaving of the attention calculations of the attention heads. In example embodiments, the

merge operation includes an arithmetic operation applied to the attention calculations of the attention heads, such as an arithmetic mean, median, min, and/or max calculation.

[3119] The attention layer outputs, for at least one token of the input sequence, a determination of attention between the token and at least one other token of the input sequence. The output of the attention calculation may include a vector that indicates, for at least one token of the input sequence, the determinations of attention between the token and a set of other tokens of the input sequence. The output of the attention calculation may include a set of vectors that indicate, for respective tokens of the input sequence, the determinations of attention between the respective token and at least one other token of the input sequence. The output of the attention calculation may indicate, for a token of a first sequence, the attention of the token to one or more tokens of a second sequence. As shown in FIG. 269, the output of the attention layer includes pairwise determinations of relatedness between pairs of tokens (e.g., each pair including a current token in an input sequence and each preceding token in the input sequence). In example embodiments, the pairwise determinations may be further processed. For example, a softmax calculation can be applied to normalize the pairwise attention determinations based on a desired range of output values (e.g., probability values between 0.0 and 1.0, with a 1.0 sum over all output values).

[3120] The attention layer may be trained by providing sets of training input sequences and comparing the outputs of the attention layer with expected outputs. Alternatively or additionally, the attention layer may be trained by incorporating the attention layer into a larger model (e.g., a transformer model) and adjusting the parameters of the attention layer (e.g., the parameters of the fully-connected layers) for a given training input sequence in order to adjust the output of the attention layer toward a desired output for the training input sequence. As an example, in a backpropagation training process, the output of the attention layer is provided as input to a succeeding layer. The output of the model including the attention layer and the succeeding layer may be compared with a desired output for the training input sequence. Based on this comparison, adjustments of the output of the succeeding layer (e.g., based on an error calculation) may inform a determination of desired adjustments of the input of the succeeding layer, which correspond to adjustments of the output of the attention layer. The adjustments of the output may be achieved by internally adjusting the parameters of the attention layer (e.g., the weights and/or biases of the fully-connected layers shown in FIG. 269) such that the attention layer subsequently generates output for the training input sequence that more closely corresponds to the desired input for the succeeding layer. Incremental training over a set of training input sequences can cause the attention layer to generate output that corresponds to the desired output for the training input sequences. As an example, if the input sequences are sentences in a language and the desired output of the model includes the probabilities of words in the language that could follow a given set of input words, the attention layer can be incrementally adjusted to indicate the attention (e.g., relatedness) between the next word in the input sequence and the preceding words in the input sequence.

[3121] It is to be appreciated that the attention layer shown in FIG. 269 presents only one example, and that attention layers may include a variety of variations with respect to the example of FIG. 269. For example, attention layers may include, without exception, additional layers or sub-layers

that perform one or more of: normalization; randomization; regularization (e.g., dropout); one or more sparsely-connected layers; one or more additional fully-connected layers; additional masking; additional reshaping and/or merging; pooling; sampling; recurrent or reentrant features, such as gated recurrence units (GRUs), long short-term memory (LSTM) units, or the like; and/or alternative layers, such as skip layers. Alternatively or additionally, the architecture of the attention layer shown in FIG. 269 may vary in numerous respects. For example, masking may be applied to the model input instead of to the outputs of the fully-connected layers. One or more fully-connected layers may be omitted, replaced with a sparsely-connected layer, and/or provided as multiple fully-connected layers, including a sequence of two or more fully-connected layers; or the like. Model parameters (e.g., weights and biases) and/or hyperparameters (e.g., layer counts, sizes, and/or embedded calculations) may be modified and/or replaced with variant parameters and/or hyperparameters. Many such variations may be included in attention layers that are incorporated in a variety of machine learning models to process a variety of types of input sequences.

Transformer Models

[3122] In example embodiments, an artificial intelligence system, machine learning model, or the like, of any of the types disclosed herein, may comprise, integrate, link to, or include a transformer model, that is, a neural network that learns context and meaning by tracking relationships in a set of sequential data inputs. Transformer models may include one or more attention layers, including (but not limited to) the attention layer shown in FIG. 269.

[3123] FIG. 270 presents an example of a transformer model. The transformer model of FIG. 270 is based on an encoder-decoder architecture in which an encoder processes an input sequence and a decoder processes an output sequence to generate output probabilities. As a first example, the input sequence may include a sequence of words in a first language; the output sequence may include a sequence of words in a second language corresponding to a translation of the input sequence; and the output probabilities may include the probabilities of words in the second language for a particular position in the translation. As a second example, the input sequence may include a sequence of words in a language that represent a query or prompt; the output sequence may include a sequence of words in the same language that represent a response to the query or prompt; and the output probabilities may include the probabilities of words in the second language for a particular position in the response. In some cases, the output sequence includes only the tokens up to a particular position (e.g., the first n-1 tokens of the output sequence), and the output probabilities represent the probabilities of tokens in the language of the output sequence that could follow the output sequence (e.g., the nth token in the output sequence). In some cases, the output sequence includes all of the tokens except the token a particular position (e.g., all of the tokens except the nth token of the output sequence), and the output probabilities represent the probabilities of tokens in the language of the output sequence that could represent the missing token in the output sequence (e.g., the nth token in the output sequence).

[3124] The encoder receives an input sequence comprising a set of tokens. The input sequence may be padded to a given length corresponding to a configured input size for the encoder. The input sequence is processed by a position encoder to encode the positions of the respective tokens

of the input sequence. The input sequence is also processed by an embedding model to determine the embeddings of the tokens of the input sequence. The encoded positions and embeddings are used to generate an encoder model input, including a query (e.g., a position of one or more tokens in the input sequence), a set of keys (e.g., the encoded positions and embeddings for each token of the input sequence), and a set of values (e.g., additional language features of the tokens such as outputs of sentiment analysis). The set of values may be a copy of the set of keys if no additional data features are available. The encoder model input is processed by a multi-head attention layer, such as an instance of the attention layer shown in FIG. 269. The multi-head attention layer determines self-attention within the input sequence (e.g., the relatedness of a respective token of the input sequence to each other token of the input sequence). The output of the multi-head attention layer is received and processed by a layer normalization component. Additionally, a skip layer is provided that passes the encoder model input through to the layer normalization component. The layer normalization component combines the output of the multi-head attention layer with the encoder model input (e.g., via arithmetic mean, median, min, max, addition, multiplication, or the like) and normalizes the combined output to within a desired range. In example embodiments, the encoder includes a sequence of two or more instances of this combination of multi-head attention layers, skip layer, and layer normalization components. The encoder also includes a feed-forward layer (e.g., a fully-connected layer and/or a sparsely-connected layer) including a set of trainable parameters. The output of the feed-forward layer is provided to another layer normalization component, along with the output of the preceding layer normalization component via a skip layer. The encoder outputs an input sequence attention, which indicates, for each of one or more tokens of the input sequence, the relatedness of each other token of the input sequence.

[3125] The decoder features an architecture that is similar to the encoder, but that includes additional components to incorporate the input sequence attention generated by the encoder. The decoder receives an output sequence comprising a set of tokens. The output sequence may be padded to a given length corresponding to a configured input size for the decoder. The output sequence is processed by a position encoder to encode the positions of the respective tokens of the output sequence. The output sequence is also processed by an embedding model to determine the embeddings of the tokens of the output sequence. The encoded positions and embeddings are used to generate a decoder model input, including a query (e.g., a position of one or more tokens in the output sequence), a set of keys (e.g., the encoded positions and embeddings for each token of the output sequence), and a set of values (e.g., additional language features of the tokens such as outputs of sentiment analysis). The set of values may be a copy of the set of keys if no additional data features are available. The decoder model input is processed by a masked multi-head attention layer, such as an instance of the attention layer shown in FIG. 269. In addition to determining attention, the masked multi-head attention layer masks the input values of a current token of the output sequence and any tokens of the output sequence that follow the current token. The masked multi-head attention layer determines self-attention within the output sequence (e.g., the relatedness of a respective token of the output sequence to each preceding token of the output

sequence). The output of the multi-head attention layer is received and processed by a layer normalization component. Additionally, a skip layer is provided that passes the encoder model input through to the layer normalization component. The layer normalization component combines the output of the multi-head attention layer with the encoder model input (e.g., via arithmetic mean, median, min, max, addition, multiplication, or the like) and normalizes the combined output to within a desired range. In example embodiments, the encoder includes a sequence of two or more instances of this combination of multi-head attention layers, skip layer, and layer normalization components. The decoder further includes an encoder-decoder multi-head attention layer that receives both the output of the preceding layer normalization component and the input sequence attention generated by the encoder. The encoder-decoder multi-head attention layer does not determine self-attention within the output sequence, but, rather, determines the attention between the tokens of the output sequence and the corresponding tokens of the input sequence. The output of the encoder-decoder multi-head attention unit is also received and processed by a second layer normalization component. Additionally, a skip layer is provided that passes the input to the encoder-decoder multi-head attention layer through to the second layer normalization component. The second layer normalization component combines the output of the multi-head attention layer with the input to the encoder-decoder multi-head attention unit (e.g., via arithmetic mean, median, min, max, addition, multiplication, or the like) and normalizes the combined output to within a desired range. The decoder also includes a feed-forward layer (e.g., a fully-connected layer and/or a sparsely-connected layer) including a set of trainable parameters. The output of the feed-forward layer is provided to a third layer normalization component, along with the output of the preceding layer normalization component via a skip layer. The output of the decoder is processed by a fully-connected layer and a softmax normalization layer based on a cross-entropy determination.

[3126] The output of the softmax normalization layer includes a set of probabilities for each possible token of a language of the output sequence for the current token. As a first example, the input sequence may include a sequence of words in a first language; the output sequence may include a sequence of words in a second language corresponding to a translation of the input sequence, up to a current (nth) word in the translation; the output probabilities may include the probabilities of words in the second language for the nth word in the translation. As a second example, the input sequence may include a sequence of words in a language that represent a query or prompt; the output sequence may include a sequence of words in the same language that represent a response to the query or prompt, up to a current (nth) word in the response; and the output probabilities may include the probabilities of words in the language for the nth word in the response.

[3127] During training, the transformer model may be provided with a set of input sequences and complete corresponding output sequences. As a first example involving language translation, the transformer model may be provided with a training data set including a first corpus of sentences in a first language and a second corpus of sentences in a second language that respectively correspond to the sentences in the first language. As a second example involving a generative model, the transformer model may be provided with a training data set including a first corpus of queries or

prompts in a language and a second corpus of responses in the language that correspond to the respective queries or prompts. For each training data input, a pair of sentences of the first corpus and second corpus are selected. The encoder is provided with the first (input) sentence, and the model is processed to determine the first word in the second (output) sentence. In this case, the output sequence provided to the decoder is completely masked so that the decoder cannot make predictions based on the expected words in the second sentence. The word probabilities determined by the decoder are compared with the actual first word in the output sequence, and backpropagation is applied through the decoder and encoder to increase the likelihood of outputting the expected word. The backpropagation includes adjusting the parameters of the attention layers to increase the attention between the first word and related words of the input sequence. The encoder is then provided again with the first (input) sentence, and the model is processed to determine the second word in the second (output) sentence. In this case, the output sequence provided to the decoder includes the unmasked first word, but masks all words after the first word. The word probabilities determined by the decoder are compared with the actual second word in the output sequence, and backpropagation is applied through the decoder and encoder to increase the likelihood of outputting the expected word. The backpropagation includes adjusting the parameters of the attention layers to increase the attention between the second word, the known first word of the output sequence, and related words of the input sequence. In this manner, the transformer model performs an autoregressive prediction, wherein the output probability of each nth token of the output sequence is based on the input sequence, the previously predicted tokens of the output sequence, and the encoder-decoder attention therebetween. Training continues over the entirety of the first and second corpora to improve the output predictions.

[3128] In many cases, the training of the transformer model occurs in batches. For example, the previous (simplified) training example described an incremental training of the transformer model over each corresponding pair of sentences of the first and second corpora, wherein the parameters of the transformer model are adjusted via backpropagation after each instance of processing. In batch training, the input and output sequences are vectorized, as are the layers of the transformer model, such that predictions over each word of the output sequence are predicted in parallel. Backpropagation parameter adjustment is performed for each batch of the training data set, based on the outputs for all of the pairwise inputs of each batch of the training data set.

[3129] After training, the transformer model can be used to predict an output sequence based on an input sequence. First, the input sequence is processed by the encoder, while the decoder processes a null output sequence (e.g., an output sequence in which all outputs are initially nulled and/or masked by the masked multi-head attention layer). The output probability of the decoder is used to determine a first token of the output sequence. In example embodiments, the first token is chosen as the token having the highest probability. In other embodiments, the first token is chosen based on a random sampling over the output probabilities. In either case, the transformer is then applied to the same input sequence and an output sequence including only the determined first token of the output sequence, and the output of the decoder determines the second token of the output sequence. This process continues until reaching an output token cap and/or upon

determining, as the output of the decoder, an end-of-sequence token. In this manner, the transformer model is applied over the input sequence to determine, in serial and autoregressive manner, the tokens of the output sequence.

[3130] It is to be appreciated that the transformer model shown in FIG. 270 presents only one example, and that transformer models may include a variety of variations with respect to the example of FIG. 270. For example, the architecture of the encoder and/or decoder may include, without exception, additional layers or sub-layers that perform one or more of: normalization; randomization; regularization (e.g., dropout); one or more sparsely-connected layers; one or more additional fully-connected layers; additional masking; additional reshaping and/or merging; pooling; sampling; recurrent or reentrant features, such as gated recurrence units (GRUs), long short-term memory (LSTM) units, or the like; and/or alternative layers, such as skip layers. Alternatively or additionally, the architecture of the encoder and/or decoder shown in FIG. 270 may vary in numerous respects. For example, masking may be applied directly to the output sequence instead of within the multi-head attention models. One or more fully-connected layers may be omitted, replaced with a sparsely-connected layer, and/or provided as multiple fully-connected layers, including a sequence of two or more fully-connected layers; or the like. Model parameters (e.g., weights and biases) and/or hyperparameters (e.g., layer counts, sizes, and/or embedded calculations) may be modified and/or replaced with variant parameters and/or hyperparameters. Many such variations may be included in transformer models to process a variety of types of input and output sequences.

[3131] Transformer models, including the example shown in FIG. 270, may be applied in a variety of circumstances. As an example, transformer models may be trained on and/or configured to process a variety of types of input sequences and/or output sequences. Sequential data inputs and/or outputs can include a wide variety of types described herein, such as strings of text, sequences of sensor data from or about an entity, sequences of steps in a process (e.g., chemical, physical, biological, and many others) or flow (e.g., a human workflow, information technology traffic flow, physical traffic flow, sequences of user behavior (e.g., attention to content, clickstream behavior, shopping behavior (digital and real world), and many others. Any of these, and others can be provided as inputs to train a transformer model, which may be alternatively described herein as a self-attention model, a foundation model, or the like. A range of mathematical self-attention techniques can be applied to detect how data elements in sequential data mutually affect each other (such as in feed forward, feedback, and other forms of influence and dependency). In various embodiments described herein and in the documents incorporated by reference herein, a set of transformer models may be deployed for a wide range of use cases, including for predictive text applications (e.g., generating a next token of text based on a previous set of tokens, such as for intelligent agent dialog, responses to queries, and the like); for extraction of information (such as extraction of meaningful elements from sensor data, signal data, and the like, such as analog signal data from sensors on machines, wearable devices, infrastructure sensors, edge and IoT devices, and many others); for analysis of human factors, such as emotional response, sentiment, satisfaction, opinion, and the like; for summarizing data (such as providing summaries of text,

images, video, sensor data, and many other streams of data of the type collected and processed as described herein); for trend detection, prediction and forecasting (and hence also for anomaly detection, such as fraud in financial transactions), including for a wide range of trends, including health (human, animal, mental, financial, machine condition, and others), performance (wellness, financial, physical, and many others), and many others; for recognition of entities and behaviors (such as objects appearing in video or image data, objects captured in LIDAR and other point-cloud rendering systems, objects located by SLAM systems, and many others); for generation and execution of instructions (e.g., recipes, control instructions, rules, regulations, governance instructions, and many others); and for many other uses.

[3132] In example embodiments, an input data set, such as an analog or digital sensor data stream, a body of text, a set of images, a set of structured data (such as data from a graph database or other form of database noted herein, a sequence of blockchain or distributed ledger entries (or other ledger data, such as accounting, financial, health or other data), a set of signals (of the various types noted herein), is provided in order to train a transformer model. In example embodiments, initial training may include a step of facilitating compression of the input data, such as by constraining the size of the transformer neural network and/or its outputs, to dimensionality that is significantly smaller (or less granular, etc.) than that of the input data. By requiring the output of the constrained transformer model to match, within a required metric of fidelity, the input data, the transformer model is caused to generate an "embedding" of the input data into a more compressed, efficient format. A decoding neural network may then be trained to operate on the output of the constrained, embedding transformer model, such that it can reproduce the input data from the output of the constrained model within the required metric, thereby assuring that the data is compressed without losing critical meaning.

[3133] Once the embedding transformer model is so trained, the decoding neural network can be removed and replaced by one or more of a set of use-case driven decoding models, each of which is trained to operate on the output of the embedding model to produce a target outcome, such as performing any of the use cases noted above to a satisfactory degree. These use-case decoding models can be fine-tuned iteratively over time with feedback from users, outcomes, or the like. Thus, a trained embedding foundation/transformer model, once created, can be used across many different use cases that may benefit from understanding the meaning of the input data set.

[3134] In example embodiments, one type of use-case decoder can be trained to allow the embedding transformer model to operate on lower quality data than was originally supplied to train the model. To accomplish this, both low quality and high quality data (such as high granularity sensor data and low granularity sensor data, or high dimensionality signal data and low dimensionality signal data, or noisy acoustic data and filtered acoustic data, or the like) can be simultaneously fed to a pair or more of instances of the trained, embedding transformer model, and a decoder for the instance of low quality data can be trained to generate an output that matches, within a metric of fidelity, the output of the instance of the embedding transformer model that is fed the high quality data. As an example, gap-free analog waveform data from a three-axis vibration sensor on a machine component can be captured simultaneously with less granular data

from a single- or two-axis accelerometer on the same component, and a decoder, operating on the output of the instance of the embedding transformer model that takes the single- or two-axis input, can be trained to match (within a tolerance) the output of the instance of the embedding transformer model that takes the more granular data as an input. Once created, the resulting decoder, coupled with the embedding transformer model, serves as a projection transformer model, effectively projecting lower quality data into higher quality data, which can then be used by other decoders to enable use cases. This class of projecting transformer models can be applied to a wide range of use cases where high quality data can be obtained during a training phase (often at higher expense), but lower quality data can be used as an input during a deployment phase (such as where lower quality data is more widely or cheaply available, such as in the case of vibration data noted above). Among other things, these projecting transformer models allow powerful, real-time, low latency use cases for AI even when input data is sparse, noisy, of low dimensionality, or the like.

[3135] In example embodiments, feedback from various decoder models can be used to improve instances of an embedding foundational or transformer model. In example embodiments, a set of transformer models, a set of decoders, or both, can be arranged in a workflow, which may be directed/acyclical or with processing loops, to create higher-level use cases that benefit from multiple applications of AI. For example, one model may be used to classify a condition, another used to generate a recommendation, and another used to generate a control instruction, among a huge range of possible embodiments. This may include serial, parallel, iterative, feed forward, feedback and other configurations.

[3136] In example embodiments, a set of models may be trained to generate instructions for configuration of other models.

[3137] In example embodiments, transformer models may be deep learning, self-learning, self-organizing, or the like, and may be used for any of the embodiments of self-learning, self-organization, or other self-referential capabilities noted throughout this disclosure or the documents incorporated by reference herein. They may also be supervised, semi-supervised, or the like. Transformer models may be coupled with, integrated with, linked to, or the like, in series, parallel or other more complex workflows, with other AI types, such as other neural network types (e.g., CNNs, RNNs, and others). For example, in example embodiments, a transformer model operating on sequential data may be coupled with a model suited to operate on non-sequential data (e.g., for pattern recognition) to achieve a use case.

[3138] In example embodiments, transformer models discover patterns in large bodies of data by application of a set of mathematical functions, optionally operating in parallel processing configurations, thereby eliminating or reducing the need for human labeling (and thereby greatly expanding the set of available data that can be used to train a model).

[3139] Self-attention may be accomplished in a transformer model by introducing a set of positional encoders that tag data elements entering and exiting a neural network and inserting a set of attention units at appropriate places in the encoding and decoding framework of an AI system. The attention units generate a mathematical map of interrelationships among data elements. In embodiments multi-headed attention units are deployed, executing a matrix of equations in parallel

to determine the interrelationships. Transformer models, using self-attention, have displayed strong capabilities to provide outputs that are consistent with how humans find patterns and meaning in data.

[3140] In example embodiments, transformer models may be embodied with very large numbers of parameters (e.g., hundreds of millions, billions, trillions, or more) operating on very large sets of parallel processors. For example, the Megatron-Turing Natural Language Generation Model by NVIDIA and Microsoft is reported to have 530 billion parameters. As noted above, from a foundational model, various use-case specific models (decoders, projections, and the like) can be purpose-built for specific applications. Accordingly, in embodiments a set of transformer models may be deployed using advanced computational techniques and/or processing architectures, such as ones that simplify or converge processors, simplify I/O, and the like. For example, 3D chipset or chiplet architectures may facilitate much higher density, faster computation, making transformer models more cost-effective. Quantum computation may also facilitate massively parallel processing in form factors that are faster, more energy efficient, or the like. Similarly, embodiments may use a tensor-engine GPU chip with a specific transformer engine, such as the NVIDIA H100 Tensor Core GPU. Another example of a transformer model is Google's switch transformer model, a trillion-parameter model that uses sparsity and a mixture-of-experts architecture to enable gains in performance and reductions in training speed.

[3141] As noted above, in embodiments smaller or more constrained transformer models may be trained to generate embeddings, particularly for very complex data sets, such as granular analog data.

[3142] In embodiments a set of transformer models may be configured to operate on structured data processing systems, such as on results from queries that are directed to a database, results of inputs directed to a set of APIs, or the like. This may facilitate better understanding of what meaning a transformer model is recognizing in a data pattern, which can be critical to ensuring quality (e.g., where a model may, due to flaws in underlying data, generate poor conclusion, such as replicating historical racial bias, missing critical balancing information, failing to understand formal logical constructs, or the like). As noted elsewhere in this disclosure and the documents incorporated herein, governance of AI in general, is a need, and the scale and complexity of transformer models likely compounds problems recognized with other neural networks, including their "black box" nature, uncertainty about input quality, and the like. Thus, governance concepts disclosed herein and in documents incorporated by reference should be understood to apply to various embodiments that use transformer models, as with other types of AI. One example is in the training of models, where models may be trained, in various disciplines, optionally similar to the educational frameworks by which humans are trained not just to sense pattern meaning, but also how to test and govern those abilities with formal reasoning and logic, mathematics, probability, and frameworks of ethics and morality.

ARTIFICIAL INTELLIGENCE AND CONVERGING TECHNOLOGY STACKS

Context

[3143] Traditional enterprise operations rely on separate, disconnected layers of technology infrastructure. Governance is largely manual, requiring significant human oversight to enforce policies, monitor compliance, and manage digital rights. Organizations struggle to maintain consistent oversight across different operational domains and often face challenges in adapting to changing regulatory requirements.

[3144] Enterprise systems typically operate with rigid, predefined offerings that lack the ability to dynamically adapt to user needs or environmental conditions. Content and services are generally standardized rather than personalized, and organizations have limited capability to customize experiences based on real-time user states, behaviors, or contextual factors.

[3145] Transaction management is predominantly handled through isolated systems with limited automation capabilities. Organizations face challenges in coordinating complex digital transactions across multiple platforms and struggle to maintain secure integration with various financial institutions and marketplaces. Digital asset management and payment systems operate with minimal intelligence, often requiring manual intervention for optimization and security enforcement.

[3146] Enterprise operations suffer from fragmented management of AI systems and technological resources. Organizations lack sophisticated capabilities for AI system generation, training, verification, and deployment. The absence of coordinated operations modules mean that AI systems were developed and deployed in isolation, without proper governance or optimization across the enterprise. Companies struggle with inefficient training data generation, inadequate model verification, and suboptimal deployment strategies that fail to leverage the full potential of their AI investments.

[3147] Enterprise networks can be characterized by static, inflexible infrastructure that often fails adapt to changing conditions. Traditional networking approaches lack intelligent routing capabilities and efficient data flow management. Organizations are unable to effectively leverage edge computing or implement dynamic network configurations, resulting in suboptimal performance and resource utilization. Network security is often reactive rather than proactive, and systems lack the ability to automatically detect and respond to potential disruptions.

[3148] Data handling is historically limited by fragmented data management approaches and isolated intelligence services. Organizations struggle to implement effective sensor and data fusion systems, resulting in incomplete or inconsistent data analysis. The absence of sophisticated machine learning systems and neural networks mean that enterprises can't fully leverage their data assets for strategic insights. Data sources remain disconnected, making it difficult to generate comprehensive intelligence or enable cross-functional data analysis.

[3149] Traditional enterprise resource optimization is characterized by manual and inefficient resource allocation processes. Organizations lack intelligent systems for optimizing the provisioning of various resources, from computational and networking resources to material and energy resources. Traditional resource management approaches are not able to effectively handle

the dynamic nature of modern enterprise operations, leading to inefficient resource utilization and increased operational costs. The absence of AI-driven optimization means that organizations often fail to balance resource allocation across different operational contexts and requirements.

[3150] The lack of intelligent integration between different technological layers created significant inefficiencies in enterprise operations. Organizations struggle to maintain coherent governance across platforms, optimize resource allocation, and deliver personalized experiences at scale. This fragmented approach limits the ability to leverage emerging technologies effectively and restricts the potential for innovation in enterprise service delivery.

[3151] These limitations in traditional enterprise systems create a clear need for a more integrated, intelligent approach to technology infrastructure.

Context

[3152] In some aspects, the techniques described herein relate to a computer-implemented system, the system including: an AI convergence system of systems.

[3153] In some aspects, the techniques described herein relate to a computer-implemented method, the method including: managing AI convergence system of systems.

[3154] In some aspects, the techniques described herein relate to an AI convergence system of systems substantially as shown and described.

[3155] In some aspects, the techniques described herein relate to a method for providing AI convergence system of systems substantially as shown and described.

[3156] In some aspects, the techniques described herein relate to a computer-implemented method substantially as hereinbefore described with reference to any of the examples and/or to the attached drawings.

[3157] In some aspects, the techniques described herein relate to a computing system including one or more processors and one or more memories configured to perform operations substantially as hereinbefore described with reference to any of the examples and/or to the attached drawings.

[3158] In some aspects, the techniques described herein relate to a computer program product residing on a computer readable storage medium having a plurality of instructions stored thereon which, when executed across one or more processors, causes at least a portion of the one or more processors to perform operations substantially as hereinbefore described with reference to any of the examples and/or to the attached drawings.

[3159] In some aspects, the techniques described herein relate to a device configured substantially as hereinbefore described with reference to any of the examples and/or to the attached drawings.

AI Convergence System of Systems

[3160] Referring to FIG. 271, an AI convergence system of systems 31000 is illustrated in accordance with some embodiments. The AI convergence system of systems 31000 represents a sophisticated multi-layered technological framework where platforms are enhanced through the deep integration of AI capabilities with other advanced technologies. This convergence fundamentally transforms how enterprise entities operate by enabling intelligent automation, data-driven decision making, and adaptive responses across all system layers. The system architecture

is designed to support complex interactions between different technological components while maintaining operational efficiency and security across the enterprise environment.

[3161] The governance layer 32000 implements comprehensive automated governance and policy enforcement through specialized governance modules 32002. This layer utilizes generative AI technology to ensure proper oversight and compliance while adapting to changing regulatory requirements. The governance modules can automatically implement and enforce policies, manage digital rights infrastructure, monitor content and data flows, and ensure compliance with relevant regulations, contracts, and licensing requirements. The layer also incorporates digital twins and reporting automation capabilities to provide real-time oversight and transparency.

[3162] The enterprise layer 33000 supports various enterprise functions by integrating enterprise management and control platforms with digital infrastructure. This layer enables automated transactional processes, manages listings and user profiles, provides analytics and insights for strategic decision-making, and enforces security and compliance protocols through sophisticated permissions systems.

[3163] The offering layer 34000 leverages an intelligent systems architecture to create, manage and enable system offerings through specialized offering modules 34002. This layer employs advanced content generation intelligence, personalization capabilities, and domain-specific knowledge processing to deliver highly customized experiences. The layer can generate and adapt content based on user profiles, behaviors, current states, environmental factors, and specific requests, while maintaining coherence and quality across multiple formats and domains.

[3164] The transactions layer 35000 provides sophisticated transaction management capabilities through a set of hardware and software-defined transaction modules. This layer implements AI models for automating, optimizing, and executing digital transactions, including digital payments and asset management. It incorporates secure integration frameworks for interfacing with financial institutions, digital marketplaces, and blockchain networks, while maintaining robust authentication and security measures.

[3165] The operations layer 36000 orchestrates the comprehensive management and execution of AI systems through specialized operations modules. These modules handle critical functions including AI system generation, training data set creation, model verification, deployment optimization, and governance oversight. The layer implements sophisticated capabilities for training AI systems using supervised, unsupervised, and reinforcement learning techniques while maintaining performance metrics and simulated data validation. Through its orchestration capabilities, this layer enables dynamic resource allocation, load balancing, and automated governance of AI workloads across the enterprise ecosystem.

[3166] The network layer 37000 delivers advanced networking capabilities through intelligent adaptation to changing conditions and requirements. This layer facilitates communications among smart network-capable systems while leveraging AI to dynamically adapt network configurations and optimize data flow. It integrates edge computing with cloud services to enable data processing and workflow orchestration closer to data sources, reducing latency and enhancing real-time decision-making capabilities. The layer employs AI algorithms for intelligent routing of various

data types through optimal network paths across cellular, WiFi, ORAN, Bluetooth and other protocols.

[3167] The data layer 38000 provides comprehensive data management and intelligence services through a sophisticated system of systems architecture. This layer implements machine learning systems, artificial intelligence systems, and neural networks configured to operate on fused data from diverse sources including sensors, wearables, social media, crowdsourced data, websites, APIs, edge devices, industrial systems, and enterprise platforms. It enables real-time monitoring and analysis while maintaining data sovereignty and security through federated learning techniques.

[3168] The resource layer 39000 enables automated and optimized provisioning of resources across various operational environments through AI-driven capabilities. This layer implements machine learning systems, artificial intelligence systems, and neural networks to optimize allocation of computational, networking, material, and energy resources. It provides dynamic resource management that can adapt to the specific demands of diverse environments like industrial settings, transportation infrastructure, energy facilities, and datacenters while maintaining operational efficiency.

[3169] This interconnected architecture creates a comprehensive technology stack that enables multiple scenarios and use cases across various enterprise contexts. The convergence of AI capabilities throughout these layers establishes an intelligent, adaptive system that can respond to diverse operational requirements while maintaining security, efficiency, and compliance. The system's modular design allows for continuous evolution and enhancement of capabilities across all layers, ensuring sustained value delivery to enterprise operations.

Governance Layer

System Overview

[3170] An example system such as a multi-layered intelligent system (e.g., AI convergence system of systems 31000) may provide automated governance (e.g., using governance layer 32000 including governance modules 32002) and execution of transactions. In example embodiments, this system may use generative AI technology to provide users and/or organizations with offerings and content that may be customized and personalized based on a profile, behavior, current state (physiological, emotional, or other), environment, location, content, and/or request. For example, a user may use this system to request a particular type of experience with a preferred theme (e.g., including music, video, characters, etc.) or a particular type of organization or presentation of content may be generated to fit with the user's style, mood and current environment.

Governance Layer Architecture

[3171] In example embodiments, the system (e.g., AI convergence system of systems 31000) may include systems for automated governance (e.g., using governance layer 32000 including governance modules 32002) and execution of transactions (e.g., using transactions layer 35000 that may include transactions modules 35002). For example, a rights infrastructure system may be included in the AI convergence system of systems 31000 (e.g., specifically in the governance layer 32000).

[3172] In some example embodiments, the system (e.g., system of systems such as SoS and more specifically AI convergence system of systems 31000) may include a governance layer 32000 that performs various governance tasks for the system. In example embodiments, the governance layer 32000 may be deployed to enforce various governance standards, workflows, or rules with respect to AI-driven tasks of the system. The types of governance standards, workflows, and rules that may be deployed in a specific configuration of the system may vary depending on the configuration of the system.

[3173] In example embodiments, the governance layer 32000 may include the set of governance modules 32002 that may be configured to perform a respective governance task. As discussed in the disclosure, governance may be configured for various purposes and in various manners. In some example embodiments, governance modules may be configured to ensure an entity is complying with a set of standards (e.g., internal standards, industry standards, recommended standards, or the like), regulations (e.g., jurisdictional regulations, professional regulations, industry regulations or the like), rules (e.g., internal rules, jurisdictional rules, international rules, or the like), or the like. Examples of the different types of analyses that a governance module may be configured to perform are described throughout the disclosure and may include risk analyses, security analyses, decision tree analyses, ethics analyses, failure mode and effects (FMEA) analyses, hazard analyses, quality analyses, safety analyses, regulatory analyses, legal analyses, and/or other suitable analyses.

[3174] In some example scenarios, a governance module may be configured to ensure an automated task is performed in a manner that complies with jurisdictional and/or governmental rules. In some example embodiments, a governance module may be configured to apply certain workflows or checklists after completion or before performance of an AI-driven task. For example, after a certain type of transaction is completed, a governance module may be configured to ensure that any required post transaction reconciliation, fulfillment, and/or reporting are completed properly. In another example, before a counterparty is transacted with, a governance module may be configured to initiate a "know-your-customer" workflow to determine whether the counterparty has a white-listed account, and if not, to determine whether to approve the counterparty transactor. In another example, a governance module may be configured and deployed to monitor a supply chain to ensure compliance with regulations relating to imported materials (e.g., tariffs are paid, regulated materials are reported, or the like).

[3175] In some example embodiments of the disclosure, a governance module may be configured and deployed to monitor the performance of an AI model (e.g., an LLM, a neural network, a machine learning model, and/or the like). In these example embodiments, a governance module may be configured to analyze the output of an AI model with respect to one or more governance parameters to determine whether the AI model is producing output that exhibits bias, drift, or the like, as described elsewhere in the disclosure. In some example embodiments, a governance module may halt use of an AI model that is determined to be exhibiting drift or bias. Additionally or alternatively, it may initiate a retraining of the model and/or alert a human user.

[3176] In some example embodiments, a governance module may be configured to analyze an input to a particular AI model. In some example embodiments, a governance module may be configured to analyze a prompt being provided to a large language model to ensure that the prompt is not requesting something that would violate a governance standard of an entity or the model. For example, a governance module may be configured to analyze a user prompt to ensure that the user is not asking the large language model to violate another party's intellectual property rights, use data sets that are not permitted to be used, or otherwise perform tasks that would put the enterprise at risk. In such example scenarios, the governance module may prevent the prompt from being processed by the LLM. Additionally or alternatively, analyzing input to an AI model may include analyzing the data being provided to the model to make a prediction or recommendation to determine whether the data is reliable data. Assuming the input data is reliable, the AI model may be provided with the input data. If, however, the data is deemed to be unreliable (e.g., fake data injection or data from a faulty sensor), a corrective measure may be initiated by the governance module (e.g., preventing the data from being used and/or alerting a human user).

Ownership and Rights Management

[3177] In example embodiments, AI agents may be configured for identifying and tracking offerings, content, systems, and services ownership for licensing. This may involve ownership of two or more entities for providing one service or multiple services relating to particular data or content. For example, the AI agents may navigate different pathways in governing ownership of services offered. The ownership may relate, for example, to copyright of media such as software, video, photos, music, and/or scripts or based on companies or individuals that provide at least a portion of the games and/or entertainment. With other services described, the ownership may relate, for example, to the roles of an individual or organization in providing service(s). There may be other logical rules used for determining ownership of content. The AI convergence system of system of systems 31000 may include management platforms that may use fully autonomous software, such as fully autonomous AI, to proceed through steps or a co-pilot software such as an AI co-pilot with a user approving and proceeding with at least some portions of the following process manually.

Content Analysis and Monitoring

[3178] In example embodiments, AI convergence system of systems 31000 may be configured for tracking and monitoring services or content going into a data model for identifying ownership of portions of the services across integrated system components or layers (e.g., across offering layer 34000 that may use offering modules 34002, enterprise layer 33000 that may include enterprise modules 33002, and governance layer 32000). In other examples, system 31000 may use various machine learning techniques and artificial intelligence capabilities, such as neural networks, to analyze content before the content data is inputted into a model for a source of the data (e.g., via offering layer 34000 and/or enterprise layer 33000). The source may indicate ownership of different parts of the content that may be assigned to one or more owners. In example embodiments, system 31000 may use intelligent processing (e.g., AI) to analyze the data going into the model as compared against continuously updated reference databases and knowledge models (e.g., a

constantly updated Internet model such as a "world" model of all public data) for identifying source information that may be used to find ownership. System 31000 may use the intelligent analysis to identify ownership from various types of information or source information associated with the content. Alternatively, in other example embodiments, automated data collection and analysis tools (e.g., a web crawler, spider, or search engine bot with AI) may be utilized to find ownership. This source information may be based on any portion of the software being used, hardware being used, and other systems being used such as relating to services and content.

Policy and Licensing Management

[3179] In example embodiments, system 31000 may use identified owners as a trigger to find any policies and/or licensing information (including e.g., contracts) that are associated with the content being inputted and/or the owners involved (e.g., using the governance layer 32000). The system 31000 may leverage the AI to find and/or locate the associated policies and/or licensing information by querying the continuously updated reference databases and knowledge models and/or querying publicly available web sources (e.g., using a web crawler, spider, or search engine bot). This querying process may find these policies and/or licensing information through direct or indirect associations to the owners identified and the related or associated content that is being inputted. In example embodiments, the querying process may also find policies and/or contractual information related to owners or at least partial owners of the related or associated services that may be requested by a user of system 31000 (e.g., owner or partial owner of services).

Transaction Management

[3180] In example embodiments, automated transaction and contract management systems (e.g., using smart contracts or other similar semi-autonomous contract systems) may be triggered via intelligent processing modules to provide transactions with a particular use of the content in modeling based on, at least partially, the associated policies, contracts, and/or licensing information discovered from querying. These automated contractual arrangements may be executed through the transaction processing components (e.g., at the transactions layer 35000 of the system 31000). In some examples, the smart contracts may be configured using distributed ledger technology (e.g., blockchain). Alternatively, in other example embodiments, communications may be automated for users to review/approve and/or automated between the system 31000 and identified owners (e.g., copyright owners of content or owners based on other factors such as development of content) to obtain approval and/or contractual information (e.g., based on, at least partially, the associated policies, contractual, and/or licensing information discovered) before proceeding with particular modeling.

Rights Compliance and Infringement

[3181] In example embodiments, system 31000 may compare results and/or output from content generation systems (e.g., generative models) against the reference knowledge bases for identifying potential rights infringement (e.g., copyright infringement) based on any substantive similarities (e.g., based on copyright law), trademark infringement, and/or infringing ownership of other rights associated with results outputted from generative models. In example embodiments, the results and/or output from generative models may be various types of services and/or other content data

generated as requested via a user interface. In order to compare results, for example, system 31000 may identify similarities between a generative model result and a "world" model and/or compared against a repository of known copyrighted works (e.g., all copyrighted media stored here with fair use data being excluded from this repository). System 31000 may use AI and/or modeling of various laws and/or regulations relating to services (e.g., statutes, rules, and case law as well as copyright law statutes, rules, and case law that relate to region(s) where the generative models are being used) for determining what is considered substantive similarities. System 31000 may use AI and/or modeling to determine whether any thresholds of infringement (e.g., copyright infringement or other ownership infringement) may have been crossed (e.g., based on the legal models or data such as statutes, rules, and case law relating to relevant jurisdiction(s)). The models used to run this analysis may be based on service-related laws and/or copyright laws (e.g., statutes, rules, and case law based on jurisdiction(s) impacted) and then may use case law examples to understand the thresholds between infringement and non-infringement. System 31000 may identify owners for portions of services being infringed or potentially infringed. System 31000 may use identified copyright owners as a trigger to find any policies, contractual, and/or licensing information (including associated contracts) that are associated with the content and/or the owners involved. Identifying owners along with associated policies and/or licensing information may occur using the same or similar approaches as described elsewhere in this disclosure. In example embodiments, system 31000 may run this process across one or more services such as comparing results and/or output from generative models against the continuously updated reference databases and knowledge models for identifying potential ownership infringement through using one or more services. The models used to run this analysis may be based on various service ownership-related laws (e.g., statutes, rules, and case law based on jurisdiction(s) impacted) and then may use case law examples to understand the thresholds between infringement and non-infringement. System 31000 may identify owners for portions of services being infringed or potentially infringed. System 31000 may use identified service owners as a trigger to find any policies, contractual, and/or licensing information (including associated contracts) that are associated with the services and/or the owners involved.

[3182] In example embodiments, system 31000 (e.g., using transactions layer 35000) may trigger smart contracts (or other similar semi-autonomous contracts) to provide transactions for particular use of resulting generative media going forward (e.g., personal use, display to a company, display to friends/family, display via internet social media channels, display on cable channels, display on streaming channels, etc.) and based on a portion of the generative result potentially infringing. In some examples, the smart contracts may be configured using distributed ledger technology (e.g., blockchain). Alternatively, communications may be generated for users to review/approve and/or automated between the system 31000 and owners to obtain approval and/or contractual information to proceed with particular use of resulting generated media based on a portion potentially infringing. This may also occur with use of resulting generated or suggested services that relate to one or more owners.

[3183] In example embodiments, system 31000 may relate to transactions layer 35000 and operations layer 36000 (e.g., that may use operations modules 36002) by use of automation software technologies to execute various systems.

Compensation and Ownership Distribution

[3184] In example embodiments, system 31000 (e.g., using enterprise layer 33000 and/or offering layer 34000) trains and deploys AI expert agents/co-pilots that can fill various artistic/creative roles within a service project. The AI expert agents may be configured to determine how to compensate rights holders that created content. For example, the AI expert agent may, when a generative work is created, determine who should be compensated and how based on identification of specific data used to train a model or similarity of the output to identified parties. In example embodiments, different components may have different licensing/compensation models. For example, gaming royalties may be different from music, video, characters, etc. For music, royalties may be awarded to individual rights holders, for example. In example embodiments, the compensation may be determined by the AI based on any combination of a potentially infringing portion of created content used, owners, as well as policies, contractual, and/or licensing information associated with the relevant owners. In some example embodiments, these ownership issues may result in fractional ownership of slot machines that may change or vary based on each instance user request as the content output may change and change during each of these instances. In example embodiments, similar processes may be executed for determining ownership percentages for services being provided especially where there may be two or more owners contributing to an offered service and/or content. Given the different ownership shares that may exist, in some example embodiments, tokenization technologies may be utilized to divide up the ownership as well as allow for the use to pay multiple owners automatically via portions of a token. In these examples, a virtual wallet may be utilized with tokens. Overall, this allows for a customized or personalized result or output based on each user's request from an interface. With services, this allows for a customized or personalized result or output based on each user's request regarding offerings, content, systems and services.

Prediction

[3185] In some example embodiments, the generative AI may suggest and/or automatically provide services to a user based on intelligent analysis of historical data and user preferences (e.g., using a prediction model). This allows the user to use any machine that is able to connect to a personalized custom model of the user for insights on the user's preferences based on historical data of previous usage and requests. The prediction process may be able to improve its prediction based on a variety of factors such as time usage of services, number of times service(s) is/are requested, success rate with some services, selection, time of day preferences, preferred spending, biometric data (e.g., using sensors to check various data points from users that can indicate likes/dislikes such as relating to relaxation vs. stress levels and/or improved focus), etc. In example embodiments, the generative AI may also suggest and/or automatically provide suggested services to a user based on a prediction model from historical data.

Model Selection

[3186] In example embodiments, the converging technology stack system 31000 selects a model for a particular task (e.g., appropriate processing models). For example, the system 31000 may determine the genus of model to use and within the genus, which specific model to use. System 31000 may select between, for example, a diffusion model or a large language model (LLM). Once a genus is selected, then the system 31000 may select a specific model (e.g., a specific diffusion model or specific LLM) to perform the task.

Digital Asset Management

[3187] In example embodiments, the system 31000 (e.g., via the transactions layer 35000) may provide for digital payment and asset management capabilities (e.g., using virtual wallets) that may include tokens and/or other payment options for transactions.

Embedded Policy and Governance, Policy Automation, and Compliance Automation

[3188] In example embodiments, the governance layer 32000 of system 31000 may include a set of machine learning systems, artificial intelligence systems, and/or neural networks configured to execute intelligence tasks associated with embedded policy and governance, policy automation, and/or compliance automation. These capabilities may work in conjunction or separately with system 31000 and other subsystems (e.g., rights infrastructure system) as described in the disclosure to provide comprehensive governance across the platform. The governance layer 32000 may automatically implement and enforce policies while ensuring compliance with relevant regulations, contracts, and/or licensing requirements.

Digital Twins and Reporting Automation

[3189] Building upon the system's ability to track ownership and manage rights, the governance layer 32000 may include and/or use digital twins (e.g., executive digital twins) and/or reporting automation capabilities that provide oversight and transparency. These digital twins may monitor governance processes in real-time, while automated reporting may ensure all stakeholders receive appropriate updates about policy implementation, compliance status, and/or system activities.

Safety Automation

[3190] The governance layer 32000 may include safety automation capabilities that may work across all operational domains. This safety automation may integrate with the system's existing monitoring and compliance features to ensure safe operations while maintaining regulatory compliance. The safety automation capabilities may adapt to different operational contexts, whether in industrial settings, healthcare environments, or transportation systems.

[3191] This interconnected architecture creates a comprehensive technology stack that enables multiple scenarios and use cases across various enterprise contexts. The convergence of AI capabilities throughout these layers establishes an intelligent, adaptive system that can respond to diverse operational requirements while maintaining security, efficiency, and compliance. The system's modular design allows for continuous evolution and enhancement of capabilities across all layers, ensuring sustained value delivery to enterprise operations.

Enterprise Layer

[3192] In embodiments, the techniques described herein relate to a computer-implemented system for managing a set of machine learning systems, a set of artificial intelligence systems, and/or set

of neural networks within an enterprise ecosystem, the system including: a processor; a memory storing instructions that, when executed by the processor, cause the system to: integrate a set of machine learning systems, a set of artificial intelligence systems, and/or set of neural networks with an enterprise access layer (EAL) that interfaces with a plurality of enterprise resources; automate the collection, storage, presentation, streaming, monitoring and analysis of data, content, processing protocols and the like, and related metadata by interfacing the a set of machine learning systems, a set of artificial intelligence systems, and/or set of neural networks within workflow systems of the enterprise; utilize a data services system to manage enterprise management and control platforms; implement an intelligence system to provide predictive analytics for trends and demand forecasting within the enterprise management and control platforms; enforce security and compliance through a permissions system that controls access to functions of the enterprise management and control platforms; manage digital transactions via a wallets system that interfaces with the enterprise management and control platforms; and generate reports on platform activity through a reporting system that is communicatively coupled with the enterprise management and control platforms

[3193] In embodiments, the techniques described herein relate to a computer-implemented system for facilitating actions within enterprise management and control platforms, including: a processor; a memory storing instructions that, when executed by the processor, cause the system to: integrate enterprise management and control platforms with an enterprise's digital infrastructure; automate transactional processes by interfacing the enterprise management and control platforms with a workflow system of the enterprise; manage listings, transactions, and user profiles using a data services system; provide analytics and insights for strategic decision-making within the enterprise management and control platforms through an intelligence system; enforce security and compliance protocols via a permissions system that controls access to the enterprise management and control platforms; and facilitate digital transactions through a wallets system that interfaces with the enterprise management and control platforms.

[3194] In embodiments, the EAL may be configured to interact with the platform users (and the ecosystem(s) in which they interact) in a variety of ways. For example, the EAL may be integrated or associated with one or more marketplaces or platforms such that the EAL functions as its own market or platform participant on behalf of the enterprise. By being associated with potentially numerous marketplaces or platforms, the EAL can perform complex or multi-stage actions, including but not limited to transactions, with enterprise assets (e.g., in a series or sequence of timed stages, simultaneously in a set of parallel transactions, or a combination of both).

[3195] In addition to marketplaces, the EAL may interact with platform users via third-party systems, some or all of which may be implemented as third-party services.

[3196] In embodiments, the EAL may include a number of EAL systems (also referred to as modules or EAL modules herein) that enable the functionality of the EAL. In some examples, these EAL systems may be deployed in a container that is specific to the EAL. When deployed in a container for the EAL, this containerized instance means that the EAL may include the necessary tools and computing resources to operate (i.e., host) the EAL systems without reliance on other

computing resources associated with the enterprise (e.g., computing resources such as processors and memory dedicated to the EAL). For example, the container for the EAL may include a set of one or more systems, such as software development kits, application programming interfaces (APIs), libraries, services (including microservices), applications, data stores, processors, etc. to execute the functions of the EAL systems that may enable the EAL to provide enterprise management and other functions and capabilities described throughout this disclosure. References herein to "EAL systems" should be understood to encompass any of the foregoing except where context dictates otherwise.

[3197] In some implementations, a set of the EAL systems may leverage computing resources considered to be external to the EAL (e.g., separate from computing resources that have been dedicated to the EAL, such as, in embodiments, computing resources shared with other enterprise applications or systems). In these implementations, the set of EAL systems leveraging external computing resources may be in communication with computing resources specific to the EAL. This type of arrangement may be advantageous when one or more of the EAL systems are computationally expensive and would increase the computational requirements for an entirely contained EAL, such as when one or more of the EAL systems causes the EAL to be a relatively expensive EAL deployment. For instance, an arrangement leveraging external (e.g., shared) systems may be beneficial for EAL systems that are infrequently utilized. To illustrate, a first enterprise may rarely use an EAL system, such as a reporting system. Here, instead of ensuring that the EAL has the computational capacity to support a reporting system by itself, the enterprise may configure the reporting system to be hosted by and/or supported by computing resources external to the EAL to deploy a relatively lean form of the EAL (i.e., an EAL container that does not include resources dedicated to a reporting system or that includes only limited resources dedicated to the reporting system with the capability to access additional, external resources as needed).

[3198] In some configurations, the EAL or a set of the EAL systems may leverage computing resources considered to be external to the EAL for support. An example of this support may be that the EAL or the set of EAL systems demands greater computing resources at some point in time (e.g., over a resource intensive time period) — for instance, greater may mean more computing resources than a normal or baseline operation state. In this example, for instance, an enterprise resource not dedicated to the EAL or EAL systems can assist or augment the services provided by some aspect of the EAL.

[3199] In embodiments, the deployment of the EAL may be configurable. For example, the enterprise or some associated developer can function as a type of architect for the EAL that best serves the particular enterprise. Additionally, or alternatively, the deployed location of the EAL may influence its configuration. For instance, the EAL may be embedded within an enterprise (e.g., non-dynamically) where it can be specifically configured using various module libraries, interface tools, etc. (e.g., as described in later detail). In some examples, the configuring entity is able to select what EAL systems will be included in its EAL. For instance, the enterprise selects from a menu of EAL systems. Here, when an EAL system is selected by the configuring entity, a

configuration routine may request the appropriate resources for that EAL system including SDKs, computing resources, storage space, APIs, graphical elements (e.g., graphical user interface (GUI) elements), data feeds, microservices, etc. In some implementations, in response to the request, the configuring entity can dedicate the identified resources of each selected EAL system. For instance, the configuring entity associates the dedicated resources to a containerized deployment of the EAL that includes the selected EAL systems.

[3200] In embodiments, the EAL may include a set of EAL systems. The set may include an interface system, a data services system, an intelligence system, a scoring system, a data pool system, a workflow system, a transaction system (also referred to as a wallet system or a digital wallet system), a governance system, a permissions system, a reporting system, and a digital twin system. Additionally, although particular types of EAL systems are described herein, the functionality of one or more EAL systems is not limited to only that particular EAL system but may be shared or configured to occur at another EAL system. For instance, in some configurations, some functionality of a transaction system may be performed by the data services system or functionality of the governance system may be incorporated with an intelligence system. In this respect, the EAL systems may be representative of the capabilities of the EAL more broadly. In embodiments, the set of EAL systems involved in any particular configuration of the EAL may include any of the systems described throughout this disclosure and the documents incorporated by reference herein, such as systems for counterparty discovery, opportunity mining, automated contract configuration, automated negotiation, automated crowdsourcing, automated facilitation of robotic process automation, one or more intelligent agents, automated resource optimization, resource tracking, and others.

[3201] In some embodiments, one or more of these systems may be configurable. The configurations may be done by selecting pre-defined configurations/plugins, by building customized modules, and/or by connecting to third party services that provide certain functionalities.

[3202] In some embodiments, aspects of a configured EAL may be dynamically reconfigured/augmented. In some examples, reconfiguration/augmentation may include updating certain data pool configurations, redefining certain workflows, changing scoring thresholds, or the like. Reconfiguration may be initiated autonomously (for example, the EAL periodically tests configurations of certain aspects of the EAL configuration using the digital twin simulation system and analytics system) or may be expert-driven (e.g., via interactions between an EAL "expert" and an interactive agent via a GUI of the interface system).

[3203] In embodiments, the data services system may perform data services for the EAL, which may include a data processing system and/or a data storage system. This may range from more generic data processing and data storage to specialty data processing and storage that demands specialty hardware or software. In some examples, the data services system includes a database management system to manage the data storage services provided by the data services system. In some configurations, the database management system may be able to perform management functions such as querying the data being managed, organizing data for, during, or upon ingestion,

coordinating storage sequences (e.g., chunking, blocking, sharding), cleansing the data, compressing or decompressing the data, distributing the data (including redistributing blocks of data to improve performance of storage systems), facilitating processing threads or queues, etc. In some examples, the data services system couples with other functionality of the EAL. As an example, operations of the data services system, such as data processing and/or data storage, may be dictated by decision-making or information from other EAL systems such as an intelligence system, a workflow system, a transaction system, a governance system, a permissions system, a reporting system, and/or some combination thereof.

[3204] It is appreciated that workflows may be deployed in any number of scenarios. Examples of scenarios where workflows may be deployed by an EAL include permission workflows, access workflows, data collection workflows, data pool workflows, machine learning workflows, artificial intelligence workflows, governance workflows, scoring workflows, transaction workflows, governance workflows, industry or vertical-specific workflows, enterprise-specific workflows, and other suitable workflows. It is appreciated that the example types of workflows provided above may overlap (e.g., a governance workflow may be an industry-specific and/or enterprise-specific workflow). Furthermore, some workflows may trigger one or more other workflows. For example, when a certain type of transaction is executed by a transaction system of an EAL, a transaction workflow corresponding to the type of transaction may define a series of tasks that are performed before the transaction is executed. In another example, as part of a data pool workflow that establishes a data pool that is accessible by third-parties, the data processing workflow may trigger a governance workflow that ensures that any enterprise data being added to the data pool confirms with certain data sharing rules (e.g., obfuscation of sensitive data, complying with privacy rules, scrubbing metadata, and/or the like) and may trigger a scoring workflow that scores each third-party that will access the data pool. Furthermore, EAL workflows may share a common framework for respective EAL functions and scenarios; however, individual workflows deployed with respect to respective EAL instances may vary in complexity from very basic workflow implementations (e.g., configured to execute on a user device or sensor device) to complex workflows with multiple dependencies and/or embedded "sub-workflows" (e.g., configured to execute by a central server system and/or by multiple enterprise devices).

[3205] In embodiments, a digital twin system may perform simulations of the enterprise's products and services that incorporate real-time data obtained from the various entities of the enterprise or third parties. In some of these embodiments, the digital twin system may recommend decisions to a user interacting with the enterprise digital twins.

[3206] In embodiments, an artificial intelligence module may include and/or provide access to a digital twin module. The digital twin module may encompass any of a wide range of features and capabilities described herein In embodiments, a digital twin module may be configured to provide, among other things, execution environments for and different types of digital twins, such as twins of physical environments, twins of robot operating units, logistics twins, executive digital twins, organizational digital twins, role-based digital twins, and the like.. In example embodiments, a digital twin module may be configured to generate digital twins that are requested by intelligence

clients. Further, the digital twin module may be configured with interfaces, such as APIs and the like for receiving information from external data sources. For instance, the digital twin module may receive real-time data from sensor systems of a machinery, vehicle, robot, or other device, and/or sensor systems of the physical environment in which a device operates. In embodiments, the digital twin module may receive digital twin data from other suitable data sources, such as third-party services (e.g., weather services, traffic data services, logistics systems and databases, and the like). In embodiments, the digital twin module may include digital twin data representing features or states, as well as demand entities, such as customers, merchants, stores, points-of-sale, points-of-use, and the like. The digital twin module may be integrated with or into, link to, or otherwise interact with an interface (e.g., a control tower or dashboard), for coordination of supply and demand, including coordination of automation within supply chain activities and demand management activities.

[3207] In embodiments, a digital twin module may provide access to and manage a library of digital twins. A plurality of artificial intelligence modules may access the library to perform functions, such as a simulation of actions in a given environment in response to certain stimuli.

[3208] In example embodiments, the digital twin(s) may be implemented with smart contracts, such as for digital twin transactions enabled by smart contracts (e.g., using smart contract orchestration engines).

[3209] In embodiments, an enterprise access layer system of systems (ELS) may manage and integrate EAL methods and systems, technological systems, data streams, platforms, and operational processes. Current enterprises face challenges in coordinating across different business units and technology stacks, creating a need for an intelligent enterprise access layer capable of parallel task execution and advanced analytics capabilities. The ELS, as described herein, implements a comprehensive architecture that enables parallel processing of intelligence tasks across multiple domains while maintaining data consistency and security protocols. The system's core architecture comprises processing unit(s) integrated with memory system(s) storing executable instructions, multiple machine learning systems, artificial intelligence engines, and neural network arrays, supported by data storage systems, communication interfaces, security modules and connections or associations with third party systems, platforms or operations.

Simulation

[3210] In embodiments, the ELS may execute system simulations utilizing neural network models to provide predictive analytics for business forecasting. In example embodiments, the system may continuously update simulations in real-time based on, for example, operational data, enabling dynamic scenario modeling and analysis for enhanced decision-making processes.

Digital Twins

[3211] In embodiments, the ELS may implement digital twin and management capabilities, generating Metaverse environments, such as industrial Metaverse environments, that enable real-time synchronization between physical and digital assets. In example embodiments, these digital representations may facilitate virtual testing and validation of industrial systems before physical implementation.

Data Integration/IoT

[3212] In embodiments, the ELS may incorporate big data processing and analytics capabilities, seamlessly integrating with Industrial Internet of Things platforms. In example embodiments, this integration may include real-time sensor data analysis and predictive maintenance optimization across industrial operations.

Process Optimization

[3213] In embodiments, the ELS may deliver automated industrial operational control through AI-driven process optimization systems. In example embodiments, the ELS may continuously perform real-time adjustments of manufacturing parameters while maintaining comprehensive quality control and compliance monitoring protocols.

Executive Digital Twins

[3214] In embodiments, the ELS may deploy executive digital twins for fleet management, enabling vehicle simulation and design optimization. In example embodiments, the ELS may manage complex fleet transaction processing while integrating with software-defined vehicle systems for enhanced operational control.

Energy Management/Digital Twins

[3215] In embodiments, the ELS may incorporate energy operations digital twins within a system of systems integration framework for comprehensive energy management. In example embodiments, the system may process energy transactions while performing AI-based energy optimization across the enterprise.

Machine/Deep Learning

[3216] In embodiments, the ELS may implement multiple machine learning approaches, including supervised learning algorithms for pattern recognition and unsupervised learning systems for anomaly detection. In example embodiments, the ELS may utilize reinforcement learning for optimization processes and deep learning capabilities for complex analysis tasks.

CNNs

[3217] In embodiments, the ELS may deploy specialized neural network architectures, including convolutional neural networks for image processing applications and recurrent neural networks for sequence analysis. In example embodiments, the system may utilize transformer networks for natural language processing tasks and implements graph neural networks for relationship analysis.

Process Optimization

[3218] In embodiments, the ELS may coordinate multiple AI systems through a hierarchical task distribution framework, implementing parallel processing optimization and resource allocation management. In example embodiments, the ELS may enable cross-system learning integration for enhanced operational efficiency.

Security/Authentication

[3219] In embodiments, the ELS may implement comprehensive security measures including role-based access control systems and multi-factor authentication protocols. In example embodiments, the system may maintain detailed activity monitoring and logging while enforcing robust security policies.

Security/Encryption

[3220] In embodiments, the ELS may implement end-to-end encryption protocols and secure data transmission mechanisms. In example embodiments, the system may incorporate privacy preservation techniques while maintaining compliance monitoring and reporting capabilities.

Enterprise Platform Integration

[3221] In embodiments, the ELS may integrate with existing enterprise resource planning systems, customer relationship management platforms, manufacturing execution systems, and supply chain management systems. In example embodiments, this integration may enable data flow and process coordination across the enterprise environment.

Remote Third-Party Connectivity

[3222] In embodiments, the ELS may maintain connections with cloud services, partner networks, and external data sources while enabling integration with third-party applications. In example embodiments, this connectivity may ensure data access and process coordination across the extended enterprise ecosystem.

Dynamic Resource Allocation

[3223] In embodiments, the ELS may implement horizontal and vertical scaling capabilities supported by dynamic resource allocation mechanisms. In example embodiments, the ELS may maintain load balancing capabilities while performing continuous performance monitoring to ensure optimal system operation.

Fault Tolerance/Redundancy

[3224] In embodiments, the ELS may incorporate fault tolerance mechanisms and redundancy management systems. In example embodiments, the implementation may include disaster recovery capabilities and continuous system health monitoring to ensure sustained operational reliability.

Production Optimization

[3225] In embodiments, in manufacturing operations, the ELS may enable production optimization through real-time monitoring and adjustment of manufacturing parameters. In example embodiments, the system may continuously analyze production data to automatically optimize resource allocation and scheduling, while maintaining quality control standards. For example, in an automotive manufacturing facility, the ELS may monitor assembly line operations in real-time, automatically adjusting robotic systems and process parameters to maintain optimal production efficiency and product quality.

Maintenance Scheduling

[3226] In embodiments, the ELS may implement predictive maintenance scheduling by analyzing equipment sensor data through its IoT integration capabilities. In example embodiments, this may allow manufacturing facilities to prevent unexpected downtime by identifying potential equipment failures before they occur. The ELS may coordinate maintenance activities across multiple production lines while optimizing resource utilization and minimizing operational disruptions.

[3227] In embodiments, the ELS may demonstrate integration capabilities by connecting manufacturing execution systems, supply chain management platforms, and enterprise resource planning systems. This integration may enable data flow and coordination across all operational

aspects. For instance, in a consumer electronics manufacturing operation, the system may coordinate production planning with supply chain logistics and inventory management, ensuring efficient resource utilization and timely delivery of finished products.

[3228] In embodiments, the ELS may implement a multi-layered neural network architecture that combines multiple specialized networks for different processing tasks. In example embodiments, the system may utilize convolutional neural networks specifically designed for image processing applications, while implementing recurrent neural networks for analyzing sequential data patterns. Additionally, transformer networks may handle natural language processing tasks, and graph neural networks process relationship analysis across the system. In these example embodiments, the ELS may coordinate these neural networks through a machine learning integration framework that implements multiple advanced approaches. For example, the system may utilize supervised learning algorithms for pattern recognition tasks, while simultaneously deploying unsupervised learning systems for anomaly detection, and/or vice versa. Reinforcement learning capabilities may be implemented for continuous optimization processes, while deep learning systems may handle complex analysis tasks requiring sophisticated pattern recognition. In embodiments, the ELS may manage the practical implementation of these systems through a hierarchical task distribution framework that enables efficient parallel processing optimization. The system may implement sophisticated resource allocation management protocols to ensure optimal utilization of computing resources across different AI subsystems. This framework facilitates cross-system learning integration, allowing different AI components to share and leverage insights across various operational domains.

[3229] In embodiments, the ELS may include a central processing unit that coordinates various AI and machine learning components through integrated memory systems storing executable instructions. In these embodiments, the system architecture may enable parallel processing of intelligence tasks across different domains while maintaining data consistency through synchronization protocols. The implementation may include dedicated data storage systems and communication interfaces that facilitate interaction between different neural network and machine learning components.

[3230] In embodiments, the ELS may implement real-time processing capabilities through its neural network arrays and AI engines. In example embodiments, these components may work in concert to process operational data, enabling continuous updates to simulations and forecasting models. In these example embodiments, the system may maintain real-time synchronization between physical and digital assets through its digital twin implementation, allowing for immediate feedback and adjustment of operational parameters.

[3231] In embodiments, the ELS may include neural network and machine learning systems that operate within a comprehensive security framework that implements end-to-end encryption and secure data transmission protocols. In example embodiments, the system may maintain privacy preservation techniques while enabling the AI components to process and analyze data without compromising security protocols.

[3232] In embodiments, the ELS may enable enterprise system integration through dedicated communication interfaces that connect with ERP systems, CRM platforms, manufacturing execution systems, and supply chain management systems. For external connectivity, the ELS may maintain connections with cloud services, partner networks, external data sources, and third-party applications through its communication interfaces and integration modules. In these example embodiments, the system-of-systems architecture may include communication interfaces and security modules that enable data synchronization while maintaining security protocols. This may include end-to-end encryption for data transmission and privacy preservation techniques for sensitive information exchange.

[3233] In embodiments, the ELS may include an integration framework that supports automated workflow management and cross-system data synchronization capabilities to ensure consistent data flow across connected enterprise systems. For performance optimization, the system may implement horizontal and vertical scaling capabilities along with dynamic resource allocation and load balancing to maintain efficient system integration and data exchange. In these example embodiments, the system's reliability features may include fault tolerance mechanisms and redundancy management to ensure consistent integration performance, while system health monitoring maintains stable connectivity with integrated enterprise systems.

[3234] Example embodiments include, but are not limited to the ELS having a set of machine learning systems, a set of artificial intelligence systems, and/or set of neural networks configured to 1) execute a set of intelligence tasks associated with a set of simulations, and 2) execute a set of intelligence tasks associated with forecasting, and 3) execute a set of intelligence tasks associated with a set of digital twins, and 4) execute a set of intelligence tasks associated with the Industrial Metaverse, and 5) execute a set of intelligence tasks associated with a value chain network control tower, and 6) execute a set of intelligence tasks associated with a set of value chain network mini control towers, and 7) execute a set of intelligence tasks associated with big data and/or analytics, and 8) execute a set of intelligence tasks associated with a set of executive digital twins, and 9) execute a set of intelligence tasks associated with an industrial internet of things (IIoT) platform, and 10) execute a set of intelligence tasks associated with AI-based industrial operational control, and 11) execute a set of intelligence tasks associated with an AI-driven Industrial Metaverse, and 12) execute a set of intelligence tasks associated with a set of financial executive digital twins, and 13) execute a set of intelligence tasks associated with enterprise transaction systems integration, 14) execute a set of intelligence tasks associated with an enterprise access layer, and 15) execute a set of intelligence tasks associated with AI-based enterprise transactional decision support, and 16) execute a set of intelligence tasks associated with a set of supply chains, and 17) execute a set of intelligence tasks associated with demand shaping integration, and 18) execute a set of intelligence tasks associated with a value chain network system of systems, and 19) execute a set of intelligence tasks associated with a set of executive digital twins for vehicle fleet operations, and 20) execute a set of intelligence tasks associated with a set of vehicle digital twins for design and simulation, and 21) execute a set of intelligence tasks associated with an enterprise access layer for fleet transactions, and 22) execute a set of intelligence tasks associated with software defined vehicle

management, and 23) execute a set of intelligence tasks associated with a set of executive digital twins for enterprise energy operations, and 24) execute a set of intelligence tasks associated with enterprise energy system of systems integration, and 25) execute a set of intelligence tasks associated with an enterprise access layer for energy transactions, and 26) execute a set of intelligence tasks associated with AI-based enterprise transactional decision support for energy management.

Offering Layer

General Systems

[3235] In embodiments, the system of systems 31000 includes an offering layer 34000 configured to create, manage, and enable offerings of the system of systems 31000 via a plurality of offering modules 34002. The offering layer includes an intelligent systems architecture that includes multiple sophisticated components configured to work together to deliver advanced capabilities across various domains. Referring to FIG. 272, the offering layer may include one or more of a content generation intelligence module 34004, a content personalization module 34006, a smart products intelligence framework module 34008, an AR processing module 34010, an immersive environment generation module 34012, a reality-virtuality fusion module 34014, a cross-reality integration module 34016, an audience segmentation module 34018, a domain-specific knowledge processing module 34020, a large language model (LLM) module 34022, a collaborative filtering module 34024, or an integration architecture module 34026.

Content Generation

[3236] In embodiments, the content generation intelligence module 34004 is configured to leverage natural language processing (NLP) models that enable automated creation of diverse content types. The NLP models employ advanced transformer architectures to understand context, maintain coherence, and generate human-like text across multiple domains and styles. One or more deep learning networks dedicated to image and video generation may utilize generative adversarial networks (GANs) and diffusion models to create high-fidelity visual content, incorporating style transfer capabilities and content-aware generation features. One or more transformer-based architectures for multi-modal content synthesis enable seamless integration of text, images, audio, and video, creating cohesive multi-format content. One or more quality assurance and content validation mechanisms employ specialized neural networks that evaluate generated content against predefined quality metrics, ensuring consistency, accuracy, and appropriateness of all outputs.

Content Personalization

[3237] In embodiments, the content personalization module 34006 implements sophisticated collaborative filtering neural networks that analyze user interactions and preferences across multiple dimensions. These networks utilize matrix factorization techniques enhanced with deep learning capabilities to identify complex patterns in user behavior. The user behavior analysis models incorporate both explicit and implicit feedback mechanisms, processing real-time interaction data to continuously refine user profiles. Real-time content adaptation engines employ reinforcement learning algorithms to dynamically adjust content presentation based on user responses and contextual factors. The personalization optimization algorithms utilize multi-armed

bandit approaches combined with deep learning to balance exploration and exploitation in content recommendations.

[3238] In embodiments, the smart products intelligence framework module 34008 includes product behavior prediction models that utilize recurrent neural networks to forecast usage patterns and potential issues before they arise. The usage pattern recognition systems employ unsupervised learning techniques to identify clusters of similar usage behaviors and anomaly detection to flag unusual patterns. Adaptive learning mechanisms continuously update product parameters based on real-world usage data, optimizing performance through reinforcement learning approaches. The real-time performance monitoring neural networks process sensor data streams to maintain optimal product functionality and predict maintenance needs.

Augmented Reality

[3239] In embodiments, the AR processing module 34010 includes one or more computer vision neural networks configured to utilize advanced convolutional architectures for real-time object recognition and tracking. Spatial mapping and tracking algorithms combine SLAM (Simultaneous Localization and Mapping) techniques with deep learning to create accurate environmental models. Real-time environment understanding models process multiple sensor inputs to maintain consistent AR experiences across varying conditions. The AR content placement optimization engines use sophisticated algorithms to determine optimal positioning of virtual elements within the physical space, considering both spatial and contextual factors.

Virtual Reality

[3240] In embodiments, the immersive environment generation module 34012 includes one or more models within the VR intelligence layer configured to create photorealistic virtual spaces using advanced rendering techniques combined with neural radiance fields. Physics simulation neural networks ensure realistic object behavior and interactions within virtual environments. User interaction prediction systems anticipate user actions to reduce latency and improve responsiveness. Experience optimization algorithms continuously adjust rendering parameters and content delivery based on hardware capabilities and user preferences.

Mixed Reality

[3241] In embodiments, the reality-virtuality fusion module 34014 enables seamless integration of virtual elements with physical environments through sophisticated scene understanding and rendering techniques. Environmental understanding neural networks process multiple sensor inputs to create accurate representations of physical spaces. Real-time scene analysis systems continuously update environmental models to maintain consistent mixed reality experiences. Interactive element placement algorithms optimize the positioning and behavior of virtual objects within the physical space.

Extended Reality

[3242] In embodiments, the cross-reality integration module 34016 includes one or more models configured to enable consistent experiences across different reality platforms through unified representation models. Universal XR content adaptation networks automatically optimize content for different display types and interaction methods. Experience synchronization algorithms

maintain temporal and spatial consistency across multiple users and devices. Multi-modal interaction processing systems handle various input methods including gesture, voice, and traditional controls.

Targeted Advertising

[3243] In embodiments, the audience segmentation module 34018 includes one or more neural networks configured to utilize deep learning techniques to identify and categorize user groups based on behavioral and demographic data. Ad performance prediction models employ ensemble learning approaches to forecast campaign effectiveness across different channels. Real-time bidding optimization systems use reinforcement learning to maximize advertising ROI. Campaign effectiveness learning algorithms continuously update targeting strategies based on performance data.

[3244] In embodiments, the domain-specific knowledge processing module 34020 includes one or more models configured to utilize sophisticated knowledge representation schemes combined with neural networks for efficient information retrieval and processing. Decision support neural networks combine traditional rule-based systems with modern deep learning approaches for improved accuracy. Inference engine optimization systems employ probabilistic reasoning techniques enhanced by machine learning. Rule-based learning algorithms automatically extract and update domain rules from new data.

Generative AI

[3245] In embodiments, the LLM module 34022 includes one or more models configured to incorporate transformer architectures with multiple (e.g., billions) of parameters for sophisticated text generation capabilities. Diffusion models for media creation utilize advanced noise prediction networks for high-quality image and video generation. Multi-modal synthesis networks enable coordinated generation of content across different modalities. Output quality validation systems employ specialized neural networks to ensure generated content meets quality and safety standards.

Recommendation Systems

[3246] In embodiments, the collaborative filtering neural networks module 34024 includes one or more models configured to implement hybrid approaches combining memory-based and model-based techniques for accurate recommendations. Content-based recommendation models utilize deep learning for feature extraction and similarity computation. Hybrid recommendation systems combine multiple approaches through sophisticated ensemble techniques. Real-time preference learning algorithms continuously update user models based on interaction data.

[3247] In embodiments, the integration architecture module 34026 is configured to implement components through a microservices-based deployment model, where each intelligence component operates as an independent service with clearly defined interfaces. Container orchestration manages resource allocation and scaling, while the API-first design ensures consistent interaction patterns across all components. The event-driven architecture enables real-time processing and communication between components through a sophisticated message broker system.

Industrial Systems

[3248] In embodiments, the offering layer includes an industrial metaverse intelligence framework that implements a sophisticated suite of artificial intelligence systems designed to enable immersive industrial operations in virtual environments. The framework employs digital twin neural networks that create high-fidelity virtual representations of physical industrial assets, processes, and facilities. These networks continuously synchronize the virtual and physical worlds through real-time sensor data integration and state prediction models. Advanced physics simulation engines powered by deep learning enable accurate modeling of industrial processes and equipment behavior within the virtual environment. The system utilizes computer vision networks for real-time tracking and reconstruction of physical spaces and assets into the metaverse environment. Natural language processing models enable intuitive voice-based interaction with virtual industrial systems and equipment. The framework includes multi-modal AI systems that facilitate seamless transitions between different interaction modalities including gesture control, voice commands, and traditional interfaces. Reinforcement learning algorithms optimize the placement and behavior of virtual industrial assets while maintaining operational constraints and safety parameters. The system employs federated learning techniques to enable distributed intelligence across multiple metaverse instances while preserving data sovereignty and security.

[3249] In embodiments, the offering layer includes an industrial system of systems integration framework that implements artificial intelligence systems designed to enable substantially seamless integration and coordination between diverse industrial systems. The framework utilizes graph neural networks to model complex relationships and dependencies between different industrial systems and subsystems. Advanced semantic reasoning engines powered by deep learning enable automated discovery and integration of industrial system capabilities and interfaces. The system employs natural language processing models to translate between different industrial protocols and data formats, enabling interoperability between legacy and modern systems. Reinforcement learning algorithms continuously optimize system-wide performance by dynamically adjusting integration parameters and resource allocation. The framework includes anomaly detection networks that identify potential integration issues or conflicts before they impact operations. Multi-agent AI systems coordinate complex interactions between integrated industrial systems while maintaining individual system autonomy. The system utilizes transfer learning techniques to apply integration patterns learned from one system pair to similar integration scenarios. Predictive analytics models leverage historical integration data to forecast potential issues and recommend preventive actions.

[3250] In embodiments, the offering layer includes an industrial control systems intelligence platform that includes neural networks and AI systems designed for advanced industrial process control and optimization. The platform employs deep reinforcement learning models for real-time process optimization and control parameter tuning. Recurrent neural networks analyze temporal patterns in process data to predict potential deviations and recommend corrective actions. The system includes computer vision networks for real-time monitoring of industrial processes and quality control. Advanced clustering algorithms automatically identify and categorize different operating states and conditions. The platform utilizes hybrid AI models that combine physics-based

modeling with data-driven learning to optimize control strategies. Predictive maintenance networks analyze equipment sensor data to forecast maintenance needs and prevent unplanned downtime. The system employs adaptive control algorithms that automatically adjust to changes in process conditions while maintaining safety constraints. Natural language processing models enable intuitive interaction with control systems through voice commands and natural language queries. The platform includes simulation capabilities powered by deep learning that enable testing of control strategies in virtual environments before deployment to physical systems.

Transactions Layer

[3251] In some example embodiments, a system-of-systems (SOS) (e.g., system of systems 31000) may include a transaction layer 35000 that includes a set of hardware and/or software-defined transaction modules that may interact with one another and/or other layers of the SOS to automate, optimize, execute, and/or otherwise facilitate digital transactions on behalf of an entity, such as an enterprise. In example embodiments, a transaction layer 35000 may deploys and/or leverages various types of AI models (e.g., generative AI models, neural networks, machine-learning models, and/or other types of models) to perform various tasks associated with a transaction (e.g., transaction execution, transaction fulfillment, transaction reconciliation, opportunity discovery, and/or the like).

[3252] In example embodiments, the transaction layer 35000 may interface with other layers of a SoS 31000. In these example embodiments, the transaction layer 35000 may provide a set of services to one or more layers modules of the SOS and/or request services from the other layers. For example, an enterprise layer 33000 may initiate a transaction in response to receiving an invoice from a third-party service or goods provider by requesting a transaction from the transaction layer. In a related example, the transaction layer 35000 may request approval services from a governance layer 32000 before the transaction layer 35000 executes a requested transaction on behalf of the enterprise using a digital wallet controlled by the transaction layer. Other examples relating to the transaction layer 35000 are discussed below.

[3253] FIG. 273 illustrates an example set of modules of an example transaction layer 35000. These may include API integrations modules 35002, authentication and security modules 35004, transaction execution modules 36005, transaction orchestration modules 35008, transaction discovery modules 35010, transaction fulfillment modules 35012, transaction reconciliation modules 35014, and/or other suitable modules.

[3254] In example embodiments, the transaction layer modules may can include and/or leverage cloud-based virtualized containerization capabilities and services, such as without limitation a container deployment and operation controller, such as Kubernetes or the like. Cloud-based virtualized containers may provide for data layers to be deployed close to source data, thereby potentially reducing network bandwidth consumption or the potential for network disturbances in a data workflow, especially when addressing voluminous local situations such as in industrial and commercial environments.

[3255] In example embodiments, technologies may be provided by and/or enabled by the data layer system of systems 38000 include intelligence services 38606, such as artificial intelligence,

machine learning and the like. These intelligence services 38606 may be provided by the data layer SoSs 38000, or accessed (e.g., as third-party services) via one or more. The data layer control 38560 may be provided access to the intelligence services 38606. In example embodiments, the data layer control 38560 may provide its own set of intelligence services.

Interfaces and Security

[3256] In example embodiments, the transaction layer 35000 may implement a secure integration framework for interfacing with external systems, such as financial institutions, digital marketplaces, and/or blockchain networks. In example embodiments, the secure integration framework may include one or more authentication and security module(s) 35006 that manage and facilitate secure connections with external systems. In example embodiments, an authentication and security module 35006 may implement respective security capabilities such as, but not limited to, multi-factor authentication, public key infrastructure (PKI) based certificate management, hardware security modules (HSMs) for cryptographic key storage and management, encrypting communication channels (e.g., using TLS 1.3 or higher), OAuth 2.0 and OpenID Connect for API authorization, IP address whitelisting, API key rotation mechanisms, and/or the like. It is appreciated that the authentication and security module 35006 may implement other respective security capabilities

[3257] In example embodiments, the transaction layer 35002 may implements an API integration framework that may provide standardized interfaces for connecting to different types of external systems (e.g., 3rd party data sources, financial institution platforms, digital marketplaces, blockchains, user devices, IoT devices, and/or the like). In these example embodiments, the transaction layer 35000 may include one or more API integration modules 35002 that implement various types of APIs and protocols including RESTful APIs for financial institution connections, WebSocket implementations for real-time market data streams, FIX protocol integration for trading systems, ISO 20022 message formats for financial transactions, GraphQL interfaces for flexible data querying, webhook handlers for asynchronous transaction notifications, and/or the like.

Transaction Execution

[3258] In example embodiments, the transaction layer 35000 may include one or more transaction execution modules 35006. In example embodiments, a transaction execution module 35006 may be configured to execute transactions on behalf of an entity. In these example embodiments, a transaction execution module 35006 may be configured to control one or more accounts of the entity. In some of these example embodiments, a transaction execution module 35006 may control one or more entity digital wallets, where each digital wallet may be configured to execute transactions from an entity account and/or one or more marketplaces. For example, a digital wallet may be configured to facilitate blockchain transactions on behalf of the entity using the credentials (e.g., private key and/or public key) of a blockchain account of the entity. In another example, a second digital wallet may be configured to securely interface with the virtual infrastructure of a respective financial institution (e.g., using an API of the financial institution and account credentials of a respective entity) to transfer funds from an account of the entity to another entity

(e.g., via a bank-to-bank wire, an account transfer within the bank, or the like). In another example, a digital wallet may interface with one or more digital marketplaces or exchanges (e.g., using an API of the marketplace/exchange). In these example embodiments, the digital wallet may be configured to transact for assets. Examples of assets that may be transacted for via a marketplace may include financial assets (e.g., foreign currencies, stocks, bonds, commodities, and/or the like), physical items (e.g., real estate, fungible goods, collectible items, and/or the like), resources (e.g., energy resources, computing resources, network resources, and/or the like). Examples of the types of transactions (or related actions) that may be performed by a digital wallet on a marketplace may include selling an asset of the entity, purchasing an asset, offering an asset for sale, making a bid on an asset, and/or the like.

[3259] In example embodiments, a transaction execution module 35006 may causes a digital wallet to perform a specific transaction (e.g., instructing a financial institution to transfer a specific amount from a specific account of the entity to a specific recipient/account; instructing a crypto wallet to generate and digitally sign a transfer request for a specific amount of cryptocurrency on a blockchain using the public key and/or private key corresponding to a specific account of the entity on the blockchain; instructing a digital wallet to purchase a specific amount of stock on an exchange; instructing a digital wallet to purchase an amount of a resource on a future or spot resource market; and/or the like).

[3260] In example embodiments, a transaction execution module 35006 may implement specialized digital wallet capabilities that execute different types of transactions using respective transaction channels. For example blockchain transactions, a transaction execution module 35006 may implement multi-signature wallet capabilities with configurable M-of-N schemes, transaction signing (e.g., using elliptic curve cryptography), gas optimization algorithms, smart contract interaction through Application Binary Interface (ABI) encoding/decoding, and/or the like. For traditional financial transactions, a transaction execution module 35006 may implement Automated Clearing House (ACH) transaction processing with NACHA file format support, Society for Worldwide Interbank Financial Telecommunication (SWIFT) message generation for international wire transfers, real-time payment system integration, automated reconciliation using ISO 20022 messages, and/or payment card processing through tokenization services. For example digital marketplace transactions, a transaction execution module 35006 may implement order book management systems, price discovery algorithms, automated trading interfaces, position management and risk monitoring capabilities, market making functionality with configurable spreads, and/or liquidity aggregation across multiple venues.

[3261] In example embodiments, a security module 35004 may allow a transaction execution module to interface with a respective transaction channels. In these example embodiments, when a transaction execution module 35006 receives a request to execute a transaction, the security module 35004 may authenticate the transaction request and establish a secure connection with the respective transaction channel corresponding to the transaction request. For example, when executing a blockchain transaction, the authentication and security module 35004 may retrieve the private key of the respective blockchain account from a secure credential vault, which may allow

a transaction execution module to digital sign the requested transaction that may be provided to a blockchain network to complete the transaction. In another example, when executing a bank transfer, the security module 36006 may retrieve the account credentials from the secure credential vault, verify the authenticity of the transaction request using various authentication protocols, and establish a secure connection with the bank's API. In example embodiments, a secure credential vault may store and manage authentication credentials used by the transaction execution modules 35006. The credential vault may implement hardware-backed key storage, an encrypted credential database, key rotation policies, access control lists, audit logging of credential usage, and/or secure credential injection into API calls.

[3262] In example embodiments, the transaction execution module 35006 may receive an instruction to perform a transaction from another transaction module, another layer of the SoS, and/or from a user associated with the entity. In example embodiments, the transaction execution module 35006 may receive a set of transaction parameters that specify information needed to execute the transaction, such as the account/digital wallet involved in the transaction, the type of transaction (e.g., sending funds, requesting funds, bidding on an item, and/or the like), the transaction amount, the account and/or digital wallet to use for the transaction, and/or any other suitable information. In response, the transaction execution module 35006 may generate an instruction to a selected enterprise digital wallet based on the transaction parameters. In response, the digital wallet may attempt to execute the transaction. In example embodiments, the digital wallet may return an outcome of the transaction (e.g., whether the transaction was successful or unsuccessful) as well as any other relevant metadata (e.g., timestamp, a receipt, an amount transferred, etc.). In response, the transaction execution module 35006 may provide the transaction outcome to various modules of the transaction layer and/or another layer of the SoS.

Transaction Orchestration

[3263] In example embodiments, the transaction layer 35000 may include one or more transaction orchestration modules 35008 that may be configured and trained to automate and/or manage one or more tasks of a transaction workflow on behalf of an enterprise. In example embodiments, the transaction orchestration module 35008 may be configured to determine a transaction orchestration workflow, initiate various tasks within a transaction orchestration workflow (which may be performed by other transaction modules of the transaction layer 35000), monitor outcomes of those tasks, and/or selectively initiate subsequent tasks based on the outcomes. In example embodiments, a transaction orchestration module 35008 may be configured to deploy an intelligent agent that may be configured and trained to orchestrate the tasks of a transaction on behalf of an enterprise, which may be referred to as a transaction orchestration agent. For example, a transaction orchestration agent transaction orchestration module 35008 may initiate a transaction discovery task that identifies a potential transaction and/or one or more potential counterparties for the potential transaction. In this example, the transaction orchestration agent may be granted authority to selectively initiate a next stage of a transaction workflow based on an outcome of the transaction discovery task or the transaction orchestration agent may be instructed by a user to initiate the next stage. In either scenario, the intelligent transaction orchestration agent may then initiate the

generation of a transaction. In some transaction workflows, the generation of a transaction may include determining a set of transaction parameters corresponding to the transaction. For example, if the transaction involves cryptocurrencies and/or a smart contract, the transaction parameters may include the blockchain address of the accounts and/or smart contract involved in the transaction (e.g., the recipient and/or sender of the cryptocurrency and/or the smart contract that is facilitating the blockchain transaction), an amount of cryptocurrency to be sent, a transaction request that is generated for digital signature, an object of the transaction (e.g., an identifier of the digital or real-world asset being transacted for, an identifier of a service being paid for, and/or the like), and/or other suitable transaction parameters. In an example of a real estate transaction, the transaction parameters may include an identifier of the property being transferred, purchaser information, seller information, a purchase price, and/or other supplemental information that may be used to generate documents relating to the real estate transaction and/or transfer of ownership of the property from one entity to a purchasing entity, and/or the like. In some example embodiments, the transaction orchestration agent may automate a contract negotiation with a counterparty to the transaction to determine the set of transaction parameters. Additionally or alternatively, a user of the system may provide some or all of the transaction parameters (e.g., via a user interface).

[3264] In some example embodiments and some types of transaction workflows, the user may designate (e.g., via a user interface) an enterprise account and/or digital wallet to perform the transaction. Alternatively, a transaction orchestration agent may be trained to select appropriate accounts for particular transactions. In some of these example embodiments, a transaction orchestration agent may be trained to determine which accounts to use for a particular transaction by analyzing various transaction parameters and conditions. For example, the transaction orchestration agent may analyze the marketplace and/or location (e.g., country, state, city) where the transaction will take place to determine whether specialized accounts are needed (e.g., accounts that are pre-approved or whitelisted for the location or are compatible with the marketplace). The transaction orchestration agent may also analyze whether the transaction involves blockchain or traditional financial transaction rails to select an account that may be configured for the respective transaction channel. The transaction orchestration agent may also analyze the transaction amount to determine whether the selected account has sufficient funds or whether the transaction amount exceeds predefined transaction limits for the account. The transaction orchestration agent may also analyze the purpose of the transaction (e.g., whether it is for operational expenses, investment purposes, or other business purposes) to select an account that may be designated for that type of transaction. The transaction orchestration agent may also analyze historical transaction patterns, risk parameters, compliance requirements, and/or enterprise policies to ensure the selected account aligns with enterprise guidelines and risk tolerance levels. In response to determining which account/digital wallet to use, the transaction orchestration module 35008 may then instruct a transaction execution module 35006 to generate and execute the transaction.

[3265] It is appreciated that a transaction orchestration agent may be trained to orchestrate other types of tasks within a transaction workflow. For example, in some example embodiments a transaction orchestration agent may initiate transaction discovery tasks. In this example, the

orchestration agent may be deployed to initiate a transaction discovery task (e.g., using a transaction discovery module), and in response to being provided a transaction opportunity, determine whether to proceed with the transaction. In another example, a transaction orchestration agent may initiate post-transaction tasks, such as fulfillment and reconciliation.

Transaction Discovery

[3266] In example embodiments, the transaction layer 35000 may include one or more transaction discovery modules. In example embodiments, a transaction discovery module 35010 may be configured to analyze data from disparate data sources for purposes of identifying potential transactions on behalf of an entity. For instance, the transaction discovery module 35010 may be configured and trained to discover potential in-bound transactions that align with an asset class or a strategy/transaction profile provided by a user. Additionally or alternatively, the transaction discovery module 35010 may be configured and trained to discover potential outbound transactions for an asset or class of assets in response to receiving an indicator of the asset or class of asset. It is appreciated that a transaction discovery module 35010 may be configured to discover different types of transactions that involve different types of asset classes. For example, in some example implementations, a transaction discovery module 35010 may be configured to identify potentially distressed assets that may be potential transaction targets. In another example, a transaction discovery module 35010 may monitor various marketplaces to identify potential arbitrage opportunities and/or opportunities to purchase energy, raw materials, and/or other resources that an enterprise may require in the future.

[3267] In example embodiments, the transaction discovery module 35010 may request or subscribe to data feeds and/or event notifications from the data layer--which may deploy intelligent agents that monitor a large number of data sources from which the data feeds and/or event notifications are mined and/or curated. The transaction discovery module 35010 may include one or more intelligence models (e.g., a set of machine learning models, neural networks, and/or other types of artificial intelligence models) and/or may interface with an intelligence system that may collectively process the data feeds and/or event notifications received from the data layer and may identify potential in-bound or outbound transactions based on the monitored data feeds and/or event notifications. In some of these example embodiments, the transaction discovery module 35010 may parse, structure, vectorize, and analyze the curated data feeds and/or event data. The transaction discovery module 35010 may then feed this data into its intelligence models, which may output potential transactions that may align with a requested asset class or transaction profile.

[3268] In example embodiments, the transaction discovery module 35010 may utilize data from disparate sources to identify potential transactions. For example, in the context of arbitrage opportunities on decentralized cryptocurrency exchanges, the transaction discovery module 35010 may monitor real-time price feeds and order book data from multiple decentralized exchanges to identify price discrepancies that may present arbitrage opportunities. In these example embodiments, the transaction discovery module 35010 may analyze liquidity pools, calculate potential arbitrage spreads while accounting for transaction costs (e.g., gas fees, slippage), identify optimal routing paths between exchanges, monitor pending transactions in a memory pool that may

impact arbitrage opportunities, and track historical price correlations and volatility patterns. The transaction discovery module 35010 may feed this data into its intelligence models to identify and score potential arbitrage opportunities.

[3269] In another example focused on discovery of real estate transaction opportunities, the transaction discovery module 35010 may monitor multiple data sources including property listings, foreclosure notices, tax records, satellite imagery, property condition reports, building permits, construction activity, comparable sales data, demographic data, and economic indicators. The transaction discovery module 35010 may analyze this data to identify properties that may match specific investment criteria (e.g., distressed properties, properties in developing areas, properties with potential for value appreciation). The transaction discovery module 35010 may use its intelligence models to parse and structure the raw data feeds, vectorize text and numerical data, apply natural language processing to unstructured content, identify patterns and correlations across data sources, generate opportunity scores and risk assessments, and rank potential transactions based on their alignment with enterprise investment criteria, etc. In identifying a potential property, the transaction discovery module 35010 may also determine an estimated value of the property and/or a proposed purchase price. In some example embodiments, the transaction discovery module 35010 may output the potential transaction opportunity to a transaction orchestration module 35008, which in turn may notify a decision maker regarding the opportunity, configure an offer to purchase the property, submit a bid on the property (e.g., if for sale or being auctioned), and/or the like.

Transaction Fulfillment and Reconciliation

[3270] In example embodiments, the transaction layer 35000 may include one or more transaction fulfillment modules 35012 that may be configured to manage and automate post-execution requirements of a transaction. In example embodiments, a transaction fulfillment module 35012 may be configured to monitor the status of a transaction after the transaction has been initiated by a transaction execution module 35006. For example, after a transaction execution module 35006 initiates a transfer of funds via a financial institution, the transaction fulfillment module 35012 may monitor the status of the transfer to determine whether the transfer was successful. In another example, after a transaction execution module 35006 initiates a blockchain transaction, the transaction fulfillment module 35012 may monitor the status of the blockchain transaction (e.g., by monitoring the number of block confirmations) to determine whether the transaction has been confirmed on the blockchain. In example embodiments, the transaction fulfillment module 35012 may generate notifications regarding the status of transactions and/or may trigger subsequent actions based on transaction status. For example, the transaction fulfillment module 35012 may generate a notification when a transaction is confirmed or may trigger the generation of transaction documentation (e.g., receipts, confirmations, transfer records, settlement instructions) when a transaction is completed.

[3271] In example embodiments, the transaction layer 35000 may include one or more transaction reconciliation modules 35014 that may be configured to verify and validate completed transactions across multiple systems of record. In example embodiments, a transaction reconciliation module

35014 may compare transaction records across internal and external systems to identify and flag any discrepancies. For example, a transaction reconciliation module 35014 may compare transaction records in an enterprise accounting system with transaction records from a financial institution to ensure that the records match. In another example, a transaction reconciliation module 35014 may compare blockchain transaction records with internal transaction records to verify that blockchain transactions were executed as intended. In example embodiments, the transaction reconciliation module 35014 may implement one or more machine learning models that may be configured and trained to detect patterns in transaction data, identify potential reconciliation issues, automate the classification of transaction exceptions, and/or suggest resolution actions for common reconciliation scenarios.

[3272] In example embodiments, the transaction reconciliation module 35014 may interface with enterprise accounting systems, financial reporting systems, and/or regulatory compliance systems to ensure accurate transaction recording and reporting. The transaction reconciliation module 35014 may implement standardized reconciliation protocols (e.g., ISO 20022) and may support various data formats (e.g., MT940, BAI2, SWIFT) for automated reconciliation processes. In example embodiments, the transaction reconciliation module 35014 may maintain audit trails of reconciliation activities and may generate exception reports for transactions that require manual review.

Industrial

Automated Industrial Marketplace

[3273] In example embodiments, an automated industrial marketplace system may integrate AI capabilities across multiple layers to facilitate industrial transactions within the transaction layer. The system may employ an industrial metaverse intelligence framework that implements digital twin neural networks to create high-fidelity virtual representations of physical industrial assets, processes, and facilities, enabling immersive industrial operations. For industrial applications, the automated industrial marketplace system may implement digital twin neural networks to create virtual representations of physical assets and processes. The system may monitor IoT data sources to determine transaction timing and attributes while coordinating resource allocation across industrial systems. For industrial transactions, the system may monitor IoT data sources to determine timing and attributes for market transactions, while utilizing IoT compute market transaction engines to coordinate resource allocation across industrial systems. The system may leverage graph neural networks to model complex relationships between different industrial systems and subsystems, while employing semantic reasoning engines powered by deep learning to enable automated discovery and integration of industrial system capabilities. Implementation may include integration with enterprise resource planning (ERP) systems to identify procurement needs, while utilizing robotic process automation (RPA) to streamline procurement processes based on inventory levels and predictive demand analysis. The system may employ reinforcement learning algorithms to optimize system-wide performance by dynamically adjusting integration parameters and resource allocation, while utilizing transfer learning techniques to apply integration patterns learned from one system pair to similar integration scenarios. The system may implement

automated governance through embedded policy and governance AI capabilities, ensuring continuous compliance monitoring and automated reporting for industrial transactions. The system relates to industrial asset transactions and resource allocation, utilizing distinct capabilities such as IoT data integration, digital twin modeling, and industrial system integration.

Operations Layer

[3274] In embodiments, the techniques described herein relate to a computer-implemented system for managing operations of a set of artificial intelligence systems. In embodiments, the computer-implemented system includes a processor and a memory storing instructions that, when executed by the processor, cause the system to store a set of operations, wherein each operation of the set of operations is associated with at least one artificial intelligence system of the set of artificial intelligence systems, and initiate at least one operation of the stored set of operations with at least one artificial intelligence system of the set of artificial intelligence systems that is associated with the least one operation.

[3275] In embodiments, the AI convergence system of systems 31000 includes an operations layer 36000 that causes the AI convergence system of systems 31000 to store a set of operations. For example, each operation may include a description of the operation to be performed. The description may include computer-readable instructions that are executable by one or more processors of one or more devices (*e.g.*, a server, a workstation, a mobile device, an edge device, a virtual device, or the like, or a combination thereof), which may be internal to and/or external to the AI convergence system of systems 31000. The description may include human-readable instructions that describe an operation in a natural language, such as English. The description may be provided at a high level of detail and/or a low level of abstraction (*e.g.*, a sequence of computer-executable instructions in machine language or a compilable programming language, a script of computer-interpretable expressions in an interpreted programming language, a declarative description of the operation in a declarative language such as XML, a flowchart, a mathematical description of an algorithm, or the like) that a processor may follow to perform the operation. The description may be provided at a low level of detail and/or a high level of abstraction (*e.g.*, a summary or overview of an operation that the AI convergence system of systems 31000 may interpret, design, generate, and execute). The description may include an indication of one or more artificial intelligence systems that may be associated with the operation, such as artificial intelligence systems that may perform the operation, upon which the operation may be performed, and/or that may participate in the design, development, generation, execution, monitoring, and/or refinement of the operation. The description may include an indication of one or more resources associated with the operation (*e.g.*, data sources that provide or receive data associated with the operation, storage devices or locations where data associated with the operation may be stored, and/or processing devices that may participate in the operation). The description may include an indication of one or more conditions under which an operation is to be initiated, performed, and/or stopped (*e.g.*, a schedule for initiating the operation on a periodic basis; a trigger condition that causes a device to initiate, perform, and/or stop the operation; and/or types of artificial intelligence systems that may be included and/or associated with a performance of the operation). The

description may include an indication of one or more outcomes of the operation (*e.g.*, an expected outcome when the operation is performed in certain conditions, and/or a log of outcomes of one or more previous performance of the operation). The set of operations may include one or more pregenerated operations that are available to be initiated in various contexts. The set of operations may include one or more *ad-hoc* operations that are generated for a particular instance, context, circumstance, device, artificial intelligence system, or the like. Following initiation and/or completion of an *ad-hoc* operation, the operation layer 36000 may continue storing the operation with the set of operations (*e.g.*, converting the *ad-hoc* operation into a pregenerated operation for future execution) or may discard the *ad-hoc* operation. The set of operations may include one or more operation templates to serve as a basis for other operations (*e.g.*, a set of device-specific stored operations that are based on a common device-independent operation template, or an *ad-hoc* operation that adapts an operation template for a particular instance).

[3276] In embodiments, the operations layer 36000 causes the AI convergence system of systems 31000 to initiate at least one operation of the set of operations with at least one artificial intelligence system. For example, the initiating may involve initiating execution of the operation on an artificial intelligence system (*e.g.*, a training operation that trains the artificial intelligence system, or an optimization operation that modifies the artificial intelligence system to optimize at least one feature). The initiating may involve invoking an artificial intelligence system to initiate the operation on another aspect of the AI convergence system of systems 31000, such as another artificial intelligence system, software based on a closed-form algorithm, an operating process, a device, a user, or the like. The initiating may involve invoking an artificial intelligence system to generate the operation (*e.g.*, invoking an artificial intelligence system to translate a description of an operation at a low level of detail / high level of abstraction into an executable operation at a high level of detail / low level of abstraction, such as an interpretable script or an executable binary). The initiating may involve invoking an artificial intelligence system to adapt an operation template into an *ad-hoc* operation based on a set of circumstances, which may be provided as input to the artificial intelligence system and/or may be generated by the artificial intelligence system. The initiating may involve invoking an artificial intelligence system to perform the operation, such as a robotic process automation (RPA) artificial intelligence system that is capable of performing a learned sequence of steps comprising a task. The initiating may involve invoking an artificial intelligence system to monitor a performance of the operation (*e.g.*, an artificial intelligence system that serves as a watchdog over the operation to ensure progress and to detect and address problems). The initiating may involve invoking an artificial intelligence system to validate a performance of the operation (*e.g.*, an artificial intelligence system that evaluates one or more outcomes of the operation to ensure that the operation has met one or more objectives for which the operation was initiated). The initiating may involve invoking an artificial intelligence system to record a performance of the operation (*e.g.*, an artificial intelligence system that records the performance of the outcome, documents one or more features of the performance of the operation, and/or summarizes a performance of an operation, one or more measures).

[3277] In embodiments, the operations layer 36000 associates resources of the AI convergence system of systems 31000 with a performance of an operation. For example, the operations layer 36000 may identify one or more artificial intelligence systems, devices, processors, storage locations, data sources, or the like that are to be included in a performance of an operation. The operations layer 36000 may secure the availability of the one or more artificial intelligence systems, devices, processors, storage locations, data sources, or the like to be included in the operation, such as reserving resources from the set of available resources of the AI convergence system of systems 31000 and/or executing transactions to acquire resources for the AI convergence system of systems 31000 for the performance of the operation. The operations layer 36000 may schedule the performance of the operation (*e.g.*, choosing a time to perform the operation in view of the execution of other planned and/or unplanned operations by the AI convergence system of systems 31000). The operations layer 36000 may organize the resources to cause the operation to be performable and/or performed (*e.g.*, communicating with a set of devices to notify the devices of their participation in the performance of the operation, and/or exchanging data among the devices that is involved in the operation). The operations layer 36000 may measure the allocation of resources to operations (*e.g.*, as a high-level planning feature to allocate the resources of the AI convergence system of systems 31000 to perform all of the planned and/or unplanned operations, to load-balance the operations in view of the resources of the AI convergence system of systems 31000, and to acquire additional resources as may be needed or desired to perform the operations).

[3278] In embodiments, the operations layer 36000 may include a set of operations modules 36002, each of which may be configured to perform and/or capable of performing one or more operations for the operations layer 36000. For example, each operations module 36002 may be specialized for various types of operations, such as a first operations module 36002 that generates artificial intelligence models; a second operations module 36002 that evaluates artificial intelligence models; a third operations module 36002 that deploys artificial intelligence models; and a fourth operations module 36002 that invokes artificial intelligence modules to perform various operations. The operations modules 36002 may be specialized for particular types of artificial intelligence modules (*e.g.*, a first operations module 36002 may be specialized for operations involving traditional artificial neural networks such as perceptron networks, a second operations module 36002 may be specialized for operations involving convolutional neural networks that include convolutional layers, and a third operations module 36002 may be specialized for operations involving large language models that include transformer layers). The operations modules 36002 may be specialized for particular types of artificial intelligence tasks (*e.g.*, a first operations module 36002 may be specialized for performing operations that involve data classification, a second operations module 36002 may be specialized for performing operations that involve computer vision, and a third operations module 36002 may be specialized for performing operations that involve content generation). The operations modules 36002 may be specialized for particular types of data and/or use cases (*e.g.*, a first operations module 36002 may be specialized for performing operations that involve healthcare data or healthcare-related use cases, a second operations module 36002 may be specialized for performing operations that

involve information technology data or information technology use cases, and a third operations module 36002 may be specialized for performing operations that involve industrial data and/or industrial use cases). In embodiments, the operations modules 36002 of the operations layer 36000 may interoperate to perform operations involving artificial intelligence systems (*e.g.*, two or more operations modules 36002 may each perform a portion of an operation involving an artificial intelligence system, such as a first operations module 36002 that trains the artificial intelligence system and a second operations module 36002 that evaluates an outcome of the training of the artificial intelligence system).

Operations Layer – Example Operations Modules

[3279] FIG. 274 is an illustration of an example operations layer 36000 that includes a set of operations modules 36002. In the example operations layer 36000 of FIG. 274, each operations module 36002 is provided to perform one set and/or type of operations. Each operations module 36002 may be implemented as one or more artificial intelligence systems, algorithms, devices, processors, data sets, or the like. Although the example operations layer 36000 of FIG. 274 depicts a particular factoring of operations to discrete modules, it is to be appreciated that some embodiments may include a different factoring of the operations to operations modules 36002. For example, in various embodiments, the operations modules 36002 of the operations layer 36000 may be merged, combined, omitted, factored into multiple modules, duplicated, repurposed, or the like. In various embodiments, the operations modules 36002 may be organized based on a different factoring principle, such as operations modules 36002 that each perform various sets of operations for particular types of artificial intelligence systems, and/or operations modules 36002 that each perform various sets of operations for particular use cases. In various embodiments, the operations modules 36002 may be organized for the operations layer 3600 in various ways, such as a centralized set of operations modules 36002 in a data center, a distributed set of operations modules 36002 provided in data centers in different geographic regions, and/or a decentralized set of operations modules 36002 that respectively perform some of the operations of the operations layer 36000 for various subsets of the AI convergence system of systems 31000, such as on behalf of different geographic regions and/or for different clients.

[3280] In embodiments, the operations modules 36002 include an AI system generation module 36004 that is configured and/or provided to generate and store artificial intelligence systems. For example, the AI system generation module 36004 may be configured and/or provided to generate artificial intelligence systems of different types and/or architectures. For example, the AI system generation module 36004 may generate artificial neural networks (ANNs) including various layers, each layer including a number of neurons and a set of interconnections to other neurons. The AI system generation module 36004 may determine the hyperparameters of the artificial neural networks, such as the number and types of inputs received by the artificial neural network, numbers and types of neural layers, the numbers and interconnections of neurons, the features of the neurons such as activation functions, additional layer features such as per-layer biases, and the number and types of outputs of the artificial neural network. For an artificial intelligence system including a random forest, the AI system generation module 36004 may determine the number and types of

inputs to the random forest, the number of decision trees included in the random forest, the sequence of features to be processed by each decision, a manner of combining the outputs of the decision trees, and the number and types of outputs of the random forest. For an artificial intelligence system including a reinforcement learning (RL) network, the AI system generation module 36004 may determine the number and types of inputs to the reinforcement learning network, the organization and capacity of the reinforcement learning network, an optimization policy to guide the functioning of the reinforcement learning network, and the number and types of outputs to the reinforcement learning network. In embodiments, the AI system generation module 36004 generates and stores artificial intelligence systems in advance of a need or use of the generated artificial intelligence systems, such as a pregenerated set of foundation artificial intelligence systems that are available for future tasks. In embodiments, the AI system generation module 36004 generates and stores artificial intelligence systems in response to a determination of a need for an artificial intelligence system to perform a task, such as determination that a software architecture in development could include an artificial intelligence system to perform a particular task such as data classification or pattern recognition. In embodiments, the AI system generation module 36004 generates and stores artificial intelligence systems in response to a request from another component of the AI convergence system of systems 31000, a processor, a device, a user, or the like, such as a request for an artificial neural network. In embodiments, the AI system generation module 36004 receives a set of hyperparameters describing a type of artificial intelligence system and generates the artificial intelligence system according to the received set of hyperparameters. In embodiments, the AI system generation module 36004 receives or determines a high-level description of an artificial intelligence system (*e.g.*, a convolutional neural network that is capable of detecting images of a selected size, resolution, bit depth, or the like), determines a set of hyperparameters of an artificial intelligence system that matches the high-level description, and generates the artificial intelligence system based on the determined set of hyperparameters. In embodiments, the AI system generation module 36004 generates a set of artificial intelligence systems that are capable of performing a task in different ways (*e.g.*, a set of convolutional neural networks configured to perform a same or similar type of computer vision task, wherein each convolutional neural network is configured to operation on images of a particular size, resolution, bit depth, or the like). In embodiments, the AI system generation module 36004 stores, in a data store of the AI convergence system of systems 31000 or elsewhere, the resources of a generated artificial intelligence system (*e.g.*, as a model file including a set of hyperparameters and parameters that comprise an artificial neural network). In embodiments, the AI system generation module 36004 provisionally reserves, in a data store of the AI convergence system of systems 31000 or elsewhere, storage capacity for the resources of an artificial intelligence system that may be generated in the future (*e.g.*, a reservation of data storage that may be sufficient for an artificial intelligence system of a particular type). In embodiments, the AI system generation module 36004 stores, in a data store of the AI convergence system of systems 31000 or elsewhere, resources for generating one or more artificial intelligence systems in the future (*e.g.*, an interpretable script, a compilable source code repository, an executable code module, or a declarative set of

hyperparameters), wherein an occurrence of a future need or desire for an artificial intelligence system may be fulfilled using the resources stored by the AI system generation module 36004. In embodiments, the AI system generation module 36004 curates a set of stored artificial intelligence systems (*e.g.*, reorganizing the stored artificial intelligence systems based on use or demand, discarding unused or underperforming artificial intelligence systems, and/or generating new artificial intelligence systems to replace less performant artificial intelligence systems).

[3281] In embodiments, the AI system generation module 36004 may generate one or more generative AI systems that are configured to generate content. For example, in response to a request to generate content of a particular type (*e.g.*, text, an image, a video, a sound, a data file, or the like) and/or having certain subjects (*e.g.*, one or more entities such as people or animals, one or more objects, one or more scenes such as a story, one or more events such as a timeline, and/or one or more interactions such as a conversation), the AI system generation module 36004 may select one or more generative AI systems that correspond to the requested content type and/or subjects. In embodiments, the AI system generation module 36004 generates a plurality of generative AI systems that produce content that is to be combined (*e.g.*, an audio generative AI system that generates audio such as music, speech, and/or sound effects, and a video generative AI system that generates still images and/or motion video, wherein the output of the combined AI systems includes a video with corresponding audio). In embodiments, the AI system generation module 36004 generates one or more AI systems to generate primary content (*e.g.*, speech) and one or more AI systems to generate supplemental content for the primary content (*e.g.*, captions and/or a transcript of the speech). In embodiments, the AI system generation module 36004 generates one or more AI systems to generate content (*e.g.*, an image) and one or more AI systems to generate metadata associated with the content (*e.g.*, a description of the content of the image). In embodiments, the AI system generation module 36004 generates one or more AI systems to generate content (*e.g.*, an image) and one or more AI systems to modify, extend, and/or otherwise further process the content (*e.g.*, one or more image filters applied to the image). In embodiments, the AI system generation module 36004 generates one or more AI systems to generate a description of content to be generated (*e.g.*, a narrative of a story) and one or more AI systems to generate content based on the description (*e.g.*, images, videos, and/or sounds that convey the story according to the narrative). In embodiments, the AI system generation module 36004 generates one or more AI systems to generate content (*e.g.*, an image) and one or more AI systems to review (*e.g.*, critique, rate, score, rank, analyze, compare, and/or verify) the content (*e.g.*, a discriminator network that identifies artifacts of the content that indicate its synthetic provenance). The AI system generation module 36004 may further configure the one or more generative AI systems to alter the content to address one or more issues indicated in the review; to generate replacement content that replaces initially generated content; and/or to be retrained to address issues identified in the review. In these and other scenarios, the AI system generation module 36004 may generate AI systems, including combinations of AI systems that operate independently and/or together, to plan, generate, supplement, review, and/or present content.

[3282] In embodiments, the AI system generation module 36004 may generate one or more AI agents that are configured to act in an autonomous manner. For example, the AI system generation module 36004 may generate an AI agent that is configured to monitor a set of conditions (*e.g.*, rules, heuristics, reflexes, or the like) and to take responsive actions based on a detected fulfillment of the conditions, even if the AI agent has not received a request, instruction, prompt, or the like to perform the responsive action. In embodiments, the AI system generation module 36004 generates an AI agent that includes a policy, heuristic, objective, goal, or the like, that identifies actions that may be taken to promote and/or maintain the policy, heuristic, objective, goal, or the like, and to initiate and/or perform the actions to promote and/or maintain the policy, heuristic, objective, goal, or the like, even if the AI agent has not received a request, instruction, prompt, or the like to perform the responsive action. The policy, heuristic, objective, goal, or the like may be specified by the AI system generation module 36004; may be specified by one or more users; and/or may be received from a data source (*e.g.*, a law, regulation, or policy defined by a government, organization, the AI convergence system of systems 31000, or the like). In embodiments, the AI system generation module 36004 configures the AI agent to learn a policy, heuristic, objective, goal, or the like based on experience (*e.g.*, by monitoring an environment such as the AI convergence system of systems 31000; by monitoring the actions of a human and the consequences of such actions (*e.g.*, robotic process automation (RPA)); by monitoring the actions of AI agents, including itself, and the consequences of such actions; or the like. In embodiments, the AI system generation module 36004 configures an AI agent to reflect on one or more policies, heuristics, objectives, goals, or the like to identify synergy (*e.g.*, actions that promote two or more policies); positive feedback loops (*e.g.*, a first policy that, if achieved, positively affects or promotes a second policy); conflicts and/or negative feedback loops (*e.g.*, a first policy that, if achieved, negatively affects or interferes with a second policy). In embodiments, the AI system generation module 36004 configures an AI agent to modify a policy, heuristic, objective, goal, or the like (*e.g.*, reducing or discarding a policy that negatively affects another policy). The AI system generation module 36004 may configure the AI agent to identify, alter, and/or operate according to priorities among two or more policies, heuristics, objectives, goals, or the like (*e.g.*, determining that a first policy is of higher priority than a second policy, and preferentially selecting actions that promote the first policy over actions that promote the second policy). In embodiments, the AI system generation module 36004 configures an AI agent to communicate about policies, heuristics, objectives, goals, or the like with AI systems or other AI agents (*e.g.*, sharing and/or comparing experiences and/or learned features) and/or to adapt policies, heuristics, objectives, goals, or the like based on information received from AI systems or other AI agents. In embodiments, the AI system generation module 36004 configures an AI agent to report on its policies, heuristics, objectives, goals, or the like, and/or about actions taken and/or not taken based on policies, heuristics, objectives, goals, or the like, to one or more AI systems (*e.g.*, supervising AI systems that occasionally review the autonomous actions of an AI agent for verification, correction, auditing, logging, or the like).

[3283] In embodiments, the operations modules 36002 include an AI system training data set generation module 36006 that is configured and/or provided to generate training data sets for training artificial intelligence systems. For example, the AI system training data set generation module 36006 may be configured and/or provided to collect data from one or more data sources for inclusion in a training data set. In embodiments, the AI system training data set generation module 36006 collects data for training data sets from one or more other modules or resources of the AI convergence system of systems 31000 (*e.g.*, from the transaction layer 35000, the network layer 37000, the data layer 38000, or the like). In embodiments, the AI system training data set generation module 36006 collects data for training data sets from other entities, facilities, or resources associated with the AI convergence system of systems 31000 (*e.g.*, from edge devices included in an industrial facility that is managed by the AI convergence system of systems 31000). In embodiments, the AI system training data set generation module 36006 collects data for training data sets from one or more external data repositories, such as news services, information repositories such as libraries, public data sources such as public databases, social networks including social media networks, communication channels, and/or the Internet. In embodiments, the AI system training data set generation module 36006 organizes and/or routes collected data into one or more training data sets. For example, the training data sets may be organized to serve the training of one or more artificial intelligence systems (*e.g.*, a first training data set for traditional artificial neural networks, a second training data set for random forests, and a third training data set for graph neural networks). The training data sets may be organized by data types (*e.g.*, a first training data set including text, a second training data set including mathematical equations or numeric measurements, and a third training data set including images). The training data sets may be organized by sources (*e.g.*, a first training data set for data collected from other modules of the AI convergence system of systems 31000, a second training data set for data collected from a local area network or a facility, a third training data set for data collected from the Internet). The training data sets may be organized by trust or reliability (*e.g.*, a first training data set of verified reliable data, a second training data set of verified unreliable data, and a third training data set of unverified data). The training data sets may be organized by provenance (*e.g.*, a first training data set of authentically generated content, a second training data set of synthetically generated content, and a third training data set of data of unknown provenance). The training data sets may be organized by use cases (*e.g.*, a first training data set including data associated with industrial manufacturing, a second training data set including data associated with healthcare, and a third training data set including data associated with a service industry). In embodiments, the AI system training data set generation module 36006 curates the training data included in each training data set (*e.g.*, normalizing, regularizing, standardizing, anonymizing, and/or cleaning the data to be included in each training data set). In embodiments, the AI system training data set generation module 36006 partitions each training data set into one or more partitions (*e.g.*, a training data set, a validation data set for periodic evaluation during training, and a test data set for use in a tiered training regimen). In embodiments, the AI system training data set generation module 36006 analyzes the training data sets for training issues (*e.g.*, underrepresentation, overrepresentation, bias, and/or

inconsistency). In embodiments, the AI system training data set generation module 36006 requests and/or directs a collection and/or creation of data to address one or more training issues associated with one or more training data sets (*e.g.*, a request or instruction to collect more data related to a particular class that is underrepresented in a training data set). In embodiments, the AI system training data set generation module 36006 generates training data sets for use with a particular training regimen, such as bootstrap aggregation ("bagging") training data sets that are partitioned in different ways for combinatorial training of an AI system In embodiments, the AI system training data set generation module 36006 stores training data sets for various training uses (*e.g.*, a first training data set for training mobile devices to perform computer vision tasks, a second training data set for training workstations to perform computer vision tasks, and a third training data set for training supercomputers to perform computer vision tasks).

[3284] In embodiments, the operations modules 36002 include an AI system training data set augmentation module 36008 that is configured and/or provided to augment training data sets for training artificial intelligence systems. For example, the AI system training data set augmentation module 36008 may be configured and/or provided to generate synthetic data to augment one or more training data sets. In embodiments, the AI system training data set augmentation module 36008 generates synthetic data in response to a request or instruction from an AI system training data set generation module 36006 (*e.g.*, based on a determination that a particular class is underrepresented in a training data set, the AI system training data set augmentation module 36008 may receive some data samples of the particular class and may generate additional synthetic data samples of the class that are similar to the received data samples of the class). In embodiments, the AI system training data set augmentation module 36008 generates synthetic data to address a particular issue with a training data set (*e.g.*, in response to a determination that a training data set of images includes too few images in low-light-level environments, the AI system training data set augmentation module 36008 may synthetically modify some existing high-light-level images to simulate low-light-level conditions). In embodiments, the AI system training data set augmentation module 36008 generates synthetic data using various augmentation techniques. For example, for training data sets including text, the AI system training data set augmentation module 36008 may generate synthetic text that modifies words and/or word ordering of authentic text; that modifies passages of text by generating replacement text; that supplements the text with generated text; and/or that combines two or more authentic passages of text to generate a synthetic text. For training data sets including images, the AI training data set augmentation module 36008 may crop, mirror, scale, shear, adjust colors, increase or decrease noise, insert or remove objects, and/or distort an authentic image to produce one or more synthetic images. For training data sets including data collected from sensors connected to industrial machines, the AI training data set augmentation module 36008 may generate synthetic data based on simulation of the sensors and/or industrial machines. The AI system training data set augmentation module 36008 may evaluate synthetic data for hallmarks of synthetic data (*e.g.*, using a discriminator network of a generative adversarial network (GAN) to determine whether synthetic data can be distinguished from authentic data) and may adapt the augmentation techniques to reduce artifacts of synthetic data. The AI system training

data set augmentation module 36008 may transmit synthetic and/or augmented data to the AI system training data set generation module 36006 for inclusion in one or more training data sets.

[3285] In embodiments, the operations modules 36002 include an AI system training module 36010 that is configured and/or provided to train artificial intelligence systems. For example, the AI system training module 36010 may be configured and/or provided to train one or more artificial intelligence systems for a particular task, such as classification, pattern recognition, object detection, and/or content generation. The AI system training module 36010 may be configured to pretrain standard models for common tasks (*e.g.*, pretraining a convolutional neural network to detect generic objects, and/or pretraining a large language model to understand common features of a natural language such as English). The AI system training module 36010 may be configured to train a generated artificial intelligence system for a particular task (*e.g.*, in response to the AI system generation module 36004 receiving a request for an artificial neural network to perform a particular classification task, the AI system generation module 36004 may generate the artificial neural network with a suitable set of hyperparameters, and the AI system training module 36010 may train the generated artificial neural network to perform the classification task). The AI system training module 36010 may be configured to train a generated artificial intelligence system using a particular training data set (*e.g.*, in response to the AI system training data set generation module 36006 receiving and/or collecting data for a particular training data set, such as a corpus of text in a particular language, the AI system training module 36010 may initiate training of a large language model train using the corpus of text). The AI system training module 36010 may be configured to adapt a pretrained model to perform a particular task (*e.g.*, upon receiving a request for a large language model that can generate a specialized type of text in a natural language, the AI system training module 36010 may identify a large language model that has previously been pretrained on the natural language and may conduct additional training of the pretrained large language model using examples of the specialized type of text). The AI system training module 36010 may use one or more types of training techniques to train AI systems, such as backpropagation, zero-shot and/or few-shot learning, supervised learning, semi-supervised learning, unsupervised learning, bootstrap aggregation ("bagging") of training data sets, reinforcement learning, genetic training, or the like. The AI system training module 36010 may evaluate a performance of an AI system during training to detect progression of the training, completion of the training, mistraining (*e.g.*, loss of performance during training), and/or training problems such as overtraining, and may adapt the training according to the evaluation (*e.g.*, ending training when the AI system exhibits performance that at least satisfies a minimum performance threshold). The AI system training module 36010 may restart training to address problems (*e.g.*, reinitiating training with a different training regimen, and/or with a different training data set, to address mistraining or overtraining of an AI system). The AI system training module 36010 may test a performance of a trained AI system (*e.g.*, validating performance of the trained models using a validation data set, and/or verifying a trained performance of an AI system using a test data set). The AI system training module 36010 may holistically allocate the resources of the AI convergence system of systems 31000 (*e.g.*, planning, scheduling, and/or adapting training of the AI systems based on performance, demand, the

availability and/or costs of compute and/or storage, or the like) to meet the need to train a set of AI systems for the set of operations of the AI convergence system of systems 31000.

[3286] In embodiments, the operations modules 36002 include an AI system verifying module 36012 that is configured and/or provided to verify generated artificial intelligence systems. For example, the AI system verifying module 36012 may be configured and/or provided to test one or more trained and/or untrained AI systems based on various criteria, such as accuracy, precision, recall, F1 score, bias, consistency, efficiency, latency, or the like. For generative AI systems such as large language models and diffusion networks, the AI system verifying module 36012 may be configured to evaluate content generated by the AI system based on features such as coherence, creativity, and variance. In embodiments, the AI system verifying module 36012 compares, scores, ranks, or otherwise considers the output of multiple AI systems (*e.g.*, to determine comparative strengths and weaknesses of the respective AI systems for performing various operations, such as generating content for particular uses). In embodiments, the AI system verifying module 36012 generates recommendations for saving, using, modifying, combining, and/or discarding various AI systems based on the verifying (*e.g.*, a first recommendation for an AI system generation module 36004 to store a first AI system that is verified as performing well; a second recommendation for an AI system training module 36010 to train or retrain a second AI system that is performing inadequately; and a third recommendation for an AI system generation module 36004 to discard and/or replace an AI system that is performing poorly). In embodiments, the AI system verifying module 36012 identifies one or more deficiencies in various AI models (*e.g.*, poor performance of the AI systems on a particular type of data, or poor performance of the AI systems in satisfying a particular task) and generates recommendations for addressing the identified deficiencies (*e.g.*, recommending an AI system training data set generation module 36006 to include additional data that is associated with an underrepresented and/or inconsistently represented class; recommending an AI system training data set augmentation module 36008 to generate certain types of augmented data, or to refrain from augmenting data in a certain way, for inclusion in training data sets; and/or recommending an AI system training module 36010 to adjust a training regimen and/or scoring model to focus on a particular issue exhibited by the AI systems).

[3287] In embodiments, the operations modules 36002 include an AI system adaptation module 36014 that is configured and/or provided to adapt a set of artificial intelligence systems to perform various operations and/or tasks. The adapting may include optimizing one or more AI systems in the performance of various operations and/or tasks, such as performing with improved accuracy, precision, recall, F1 scores, bias, consistency, efficiency, latency, or the like. Herein, the term "optimizing" may mean an improvement either to a globally optimal or best state, locally optimal or best state within a given domain, improving relative to a previous state, or the like. For example, the AI system adaptation module 36014 may be configured and/or provided to receive a request to adapt a pretrained AI system (*e.g.*, a partially trained large language model) for a specific task, device, data type, deployment, use case, or the like. For example, the AI system adaptation module 36014 may be configured to adapt an AI system that is configured to receive, process, and output a first type of data (*e.g.*, images of a first size, resolution, bit depth, or the like) so that the AI

system can process a second type of data (*e.g.*, images of a second size, resolution, bit depth, or the like). The AI system adaptation module 36014 may be configured to adapt an AI system that is configured to receive, process, and output data in a first context (*e.g.*, text associated with documents of a first type) so that the AI system can process a second type of data (*e.g.*, text associated with documents of a second type). The AI system adaptation module 36014 may be configured to adapt an AI system that is configured to receive, process, and output data in a first context (*e.g.*, data received from sensors in an industrial environment) so that the AI system can process data in a second context (*e.g.*, data received from sensors in a healthcare environment). The AI system adaptation module 36014 may be configured to adapt an AI system that is configured to receive, process, and output data on a first type of device (*e.g.*, a workstation including a graphics processing unit (GPU)) so that the AI system can process data on a second type of device (*e.g.*, an embedded device featuring only a traditional central processing unit (CPU) and/or a field-programmable gate array (FPGA)). The AI system adaptation module 36014 may be configured to adapt an AI system that is configured to receive, process, and output data to perform a first task on a data set (*e.g.*, data classification) so that the AI system can process data on a second task on the same data set (*e.g.*, pattern detection). In view of these and other types of adaptations, the AI system adaptation module 36014 may identify one or more AI systems stored by an AI system generation module 36004 and one or more adaptations of the one or more AI systems that may adapt the AI system according to a request or instruction. The adaptation may include preprocessing an input to an AI system to alter the input received by the AI system, such as adding a preprocessing denoising filter to the input of an AI system to reduce noise and improve the processing, output, and performance of the denoised data. The adaptation may include altering an AI system to receive a different type of input, such as adapting a convolutional neural network to receive and process images of a different size, resolution, color depth, or the like than the images on which the convolutional neural network was previously trained. The adaptation may include reconfiguring an architecture of an AI system, such as adding one or more architectural features (*e.g.*, a long short-term memory (LSTM) unit, a gated recurrence unit (GRU), a transformer layer, a filter layer, a pooling layer, or the like), removing one or more architectural features, replacing one or more architectural features with a different architectural feature, or the like. The adaptation may include altering an AI system to process data differently, such as appending one or more fine-tuning layers to a pretrained artificial neural network to provide additional processing for a specific set of data. The adaptation may include altering an AI system to generate different output, such as adapting an artificial neural network that processes input to select a class from a first set of classes, such that the artificial neural network performs the same or similar processing of the input to select a class from among a second, different set of classes. The adaption may include postprocessing an output generated by the AI system to alter an output of the AI system, such as adding a postprocessing denoising feature to reduce noise in content generated by a generative AI system. The adaptation may include selectively limiting a processing of input by the AI system to a subset of input for which the AI system is determined to perform well. The adaptation may include requesting a continued training or retraining of an AI system by the AI system training module

36010, comparing a performance of the additionally trained AI system with the original AI system, and determining one or more adaptations of the AI system based on the comparison. The adaptation may include altering a set of resources provided to the AI system (*e.g.*, altering an amount of memory allocated for the AI system, an amount or type of computation performed by or for the AI system, such as to reduce latency of a time-sensitive model and/or a context window size of a time series model or a large language model to improve memory and/or performance). In embodiments, the AI system adaptation module 36014 may consider a set of possible adaptations of an AI system for addressing one or more issues or problems of the AI system (*e.g.*, poor accuracy, poor consistency, poor coherence, high latency, or the like), comparing the predicted and/or determined effects of various adaptations and combinations thereof of the AI system, selecting one or more adaptations to apply to the AI system based on the comparison, and adapting the AI system based on the selection.

[3288] In embodiments, the operations modules 36002 include an AI system aggregation module 36016 that is configured and/or provided to aggregate artificial intelligence systems. The AI system aggregation module may combine two or more AI systems into a hybrid AI system that combines the heterogeneous processing capabilities of AI systems with different architectures. For example, the AI system aggregation module 36016 may combine an AI system with one or more other AI systems to produce a hybrid system that generates aggregated and improved output as compared with the individual AI systems (*e.g.*, a "boosted" ensemble of inadequately performing AI systems, from which a consensus output may be derived that presents higher-performance output than from any of the individual AI systems). The aggregation may include adding an AI system to a set of AI systems, each of which may be determined to perform well over a subset of input, wherein a particular input to the set of AI systems may be selectively processed by a particular AI system that is recognized to perform well over the subset of input that includes the particular input. In embodiments, the AI system aggregation module combines two or more AI systems into an ensemble AI system that features a collection of homogeneous or heterogeneous AI systems that together generate an output (*e.g.*, by consensus). In embodiments, the AI system aggregation module combines two or more AI systems into a sequence or chain of AI systems, wherein an output of a first AI system is received as input by a second AI system. The aggregation may include connecting at least one output of a first AI system to at least one input of a second AI system (*e.g.*, a chain-of-processing generative system in which data generated by a generative AI system is refined, expanded, reformatted, verified, analyzed, augmented, edited, or otherwise further processed by a second AI system). In some embodiments, the AI system aggregation module 36016 generates a swarm of AI systems that operate independently of one another but collectively to achieve a task, such as a set of AI agents that are configured to operate individually but that interact with other AI agents and shared resources to achieve the task. In some embodiments, the AI system aggregation module 36016 aggregates AI systems by inserting at least a portion of a first AI system into a second AI system (*e.g.*, inserting one or more layers of a classifier network between two layers of a convolutional neural network to inform further processing of the convolutional neural network based on the classification of the partially-processed input data). The aggregation may

include conditioning the processing of data by a first AI system based on the processing of the data by a second AI system (*e.g.*, partially generating content by a generative AI system, and then conditioning a completion of the generated content on an initial assessment of the partially generated content by a second AI system, such as a discriminator network). The aggregation may include combining a first AI system to select, design, and/or validate a processing of data by a combination of one or more other AI systems (*e.g.*, an expert determining AI system that determines, from a set of AI systems that are respectively experts in various domains, one or more AI systems that are to be invoked to process input data that is relevant to one or more domains). The aggregation may include attaching a first AI system as an administrator, regulator, controller, or the like of a second AI system (*e.g.*, a content moderation AI system that determines whether a generative AI system is permitted to process certain kinds of prompts and/or generate certain kinds of data, and that controls the processing of the generative AI system based on the determined permission). In embodiments, the AI system aggregation module 36016 identifies one or more aggregations of pregenerated AI systems for various contexts (*e.g.*, combinations of pregenerated AI systems that may be specialized and/or proficient at certain types of processing, such as generating certain forms of content, generating classifications above a classification threshold, and/or completing processing within a latency window). In embodiments, the AI system aggregation module 36016 generates and stores one or more descriptions of techniques for selecting and aggregating AI systems for various contexts (*e.g.*, an interpretable script for selecting and combining AI systems to generate chain-of-thought reasoning for various contexts), wherein a particular task or invocation of the AI systems may be fulfilled by applying one of the aggregation techniques generated by the AI system aggregation module 36016 to select and aggregate AI systems for the indicated context. Such aggregating may be advantageous, for example, for *ad-hoc* aggregations of AI systems for unusual, specialized, and/or previously unseen processing tasks (*e.g.*, a request to aggregate AI systems to generate a hybrid large language model for execution on a particular device and/or over a particular type of data).

[3289] In embodiments, the operations modules 36002 include an AI system deployment module 36018 that is configured and/or provided to deploy artificial intelligence systems. For example, the AI system deployment module 36018 may be configured and/or provided to deploy one or more AI systems generated by the AI system generation module 36004 to one or more devices, such as a set of workstations, servers, or mobile devices. For example, the AI system deployment module 36018 may receive a request to deploy a large language model to a device having a certain set of computational resources, such as processing cores operating at a defined clock rate, memory, and storage. The AI system deployment module 36018 may review a set of AI systems generated by the AI system generation module 36004 to identify a suitable large language model in view of the computational resources of the device. The AI system deployment module 36018 may select the large language model for deployment, determine a location of the resources of the selected large language model, package the resources of the large language model for deployment (*e.g.*, retrieving and compressing a description of the architecture of the large language model and a set of parameters of the large language model), and transmit the packaged resources for deployment on

the device. In embodiments, the AI system deployment module 36018 is configured to deploy an AI system for a particular type of task, such as a particular classification task, pattern recognition task, computer vision task, or content generation task. The AI system deployment module 36018 may select a particular AI system from the set of AI systems generated by the AI system generation module 36004 to identify an AI system that is suitable for the identified task. The AI system deployment module 36018 may package the identified AI system for deployment to the device and/or environment for the identified task, and may transmit the packaged AI system to the device and/or environment. In embodiments, the AI system deployment module 36018 is configured to deploy an AI system to include a particular feature set, such as a large language model that supports a particular natural language and features a particular context window size, and may deploy a pregenerated or adapted AI system to fulfill the deployment that features the indicated feature set. In embodiments, the AI system deployment module 36018 is configured to deploy an AI system for execution in a particular execution environment, such as within a defined memory space, using a particular AI framework such as TensorFlow or Torch, and/or using a particular set of hardware that is available to the device, such as a particular graphics processing unit (GPU), neural processing unit (NPU), tensor processing unit (TPU), or the like. In embodiments, the AI system deployment module 36018 is configured to adapt an AI system for deployment to a particular device and/or context (*e.g.*, identifying an AI system that is suitable for a particular task but that is based on a TensorFlow environment, adapting the AI system for execution in a Torch environment, and transmitting a packaged version of the Torch-based AI system to a device featuring a Torch execution environment).To achieve the adapting for the deployment, the AI system deployment module 36018 may transmit a request to the AI system adaptation module 36014 to adapt a stored AI system in view of the details of a particular deployment, may receive an adapted AI system from the AI system adaptation module 36014 based on the request, and may package and deploy the adapted AI system to satisfy the deployment. In embodiments, the AI system deployment module 36018 may transmit the resources of one or more AI systems to one or more devices that are intended to execute the AI systems. In embodiments, the AI system deployment module 36018 may transmit the resources of an AI system to two or more devices for redundant and/or distributed execution of the AI system (*e.g.*, as a primary system and a failover system, as two or more devices that provide scalable capacity for execution of the AI systems, as a distributed system in which a first system executes a first AI system that transmits an output to a second device for input to a second AI system, and/or as a federated system in which respective devices execute a portion of the AI system). In embodiments, the AI system deployment module 36018 may deploy an AI system within the AI convergence system of systems 31000, such as by reserving resources (*e.g.*, processing capacity, storage, one or more dedicated devices, or the like) for the execution of an AI system to serve a particular context, and/or by transmitting the resources of an AI system (*e.g.*, the parameters of an artificial neural network) to one or more devices within the AI system deployment module 36018. In embodiments, the AI system deployment module 36018 may deploy an AI system by transmitting an interface to the AI system to one or more devices, such as end-user devices that are configured to collect input, transmit input through the interface to one or more AI

systems executed on other devices, and/or receive and present output received from one or more AI systems. For example, the AI system deployment module 36018 may retrieve, generate, and/or transmit an application programming interface (API) to the one or more devices in order to interconnect the one or more devices with the AI system. The AI system deployment module 36018 may reserve one or more communication endpoints to enable one or more devices to communicate with the AI system. The AI system deployment module 36018 may configure permissions and/or security tokens to enable one or more external devices to communicate securely with one or more AI systems.

[3290] In embodiments, the operations modules 36002 include an AI system invocation module 36020 that is configured and/or provided to invoke artificial intelligence systems to perform various operations and/or tasks. For example, the AI system invocation module 36020 may be configured and/or provided to receive one or more requests to perform a task. The request may be initiated by another component of the AI convergence system of systems 31000; by a device, software process, or entity that is served by the AI convergence system of systems 31000; and/or by an external device, software process, or entity. The request may be spontaneously initiated by the AI system invocation module 36020 (*e.g.*, based on a set of conditions monitored by the AI system invocation module 36020, wherein fulfillment of a condition instructs the AI system invocation module 36020 to invoke an AI system). The request may indicate one or more AI systems that are to be invoked to execute the request. Alternatively or additionally, the request may indicate one or more tasks, and the AI system invocation module 36020 may determine one or more AI systems that may be invoked to fulfill the task. In some cases, the request may specify a request that includes one or more non-AI systems (*e.g.*, a hardware component or a traditional software component, such as a graphics rendering or simulation system), and the AI system invocation module 36020 may identify one or more AI systems that may be invoked to perform the task in place of the one or more non-AI systems. In embodiments, the AI system invocation module 36020 verifies a permission and/or security of a request to invoke an AI system before initiating the invocation of the AI system. In embodiments, the AI system invocation module 36020 initiates the invocation of the AI system in response to the request (*e.g.*, by instantiating the AI system, by transmitting input to the AI system, and/or by securing the resources for the invocation of the AI system). In embodiments, the AI system invocation module 36020 notifies a requester of the status of a request to invoke an AI system (*e.g.*, indicating a status, progress, estimated completion time, and/or outcomes of the processing of the request by the AI systems). In embodiments, the AI system invocation module 36020 mediates communication between the AI system and the requester (*e.g.*, by transmitting the request and/or input received from the requester and/or another source to the AI system to invoke the AI system for the request, and/or by transmitting one or more outputs of the AI system to the requester and/or another destination to fulfill the request). In embodiments, the AI system invocation module 36020 translates the request and/or one or more inputs received from the requester and/or another source for intake and processing by the AI system, and/or translates one or more output of the AI system for transmission to the requester and/or another destination).

[3291] In embodiments, the AI system invocation module 36020 may invoke one or more generative AI systems that are configured to generate content in response to a request. For example, in response to a request to generate content of a particular type (*e.g.*, text, an image, a video, a sound, a data file, or the like) and/or having certain subjects (*e.g.*, one or more entities such as people or animals, one or more objects, one or more scenes such as a story, one or more events such as a timeline, and/or one or more interactions such as a conversation), the AI system invocation module 36020 may select and invoke one or more generative AI systems that correspond to the requested content type and/or subjects. In embodiments, the AI system invocation module 36020 invokes a plurality of generative AI systems that produce content that is to be combined (*e.g.*, an audio generative AI system that generates audio such as music, speech, and/or sound effects, and a video generative AI system that generates still images and/or motion video, wherein the output of the combined AI systems includes a video with corresponding audio). In embodiments, the AI system invocation module 36020 invokes one or more AI systems to generate primary content (*e.g.*, speech) and one or more AI systems to generate supplemental content for the primary content (*e.g.*, captions and/or a transcript of the speech). In embodiments, the AI system invocation module 36020 invokes one or more AI systems to generate content (*e.g.*, an image) and one or more AI systems to generate metadata associated with the content (*e.g.*, a description of the content of the image). In embodiments, the AI system invocation module 36020 invokes one or more AI systems to generate content (*e.g.*, an image) and one or more AI systems to modify, extend, and/or otherwise further process the content (*e.g.*, one or more image filters applied to the image). In embodiments, the AI system invocation module 36020 invokes one or more AI systems to generate a description of content to be generated (*e.g.*, a narrative of a story) and one or more AI systems to generate content based on the description (*e.g.*, images, videos, and/or sounds that convey the story according to the narrative). In embodiments, the AI system invocation module 36020 invokes one or more AI systems to generate content (*e.g.*, an image) and one or more AI systems to review (*e.g.*, critique, rate, score, rank, analyze, compare, and/or verify) the content (*e.g.*, a discriminator network that identifies artifacts of the content that indicate its synthetic provenance). Based on the review, the AI system invocation module 36020 may invoke one or more AI systems to alter the generated content to address one or more issues indicated in in the review; to generate replacement content that replaces initially generated content; and/or to retrain the generative AI systems to address issues identified in the review. In these and other scenarios, the AI system invocation module 36020 may invoke AI systems, including combinations of AI systems that operate independently and/or together, to plan, generate, supplement, review, and/or present content.

[3292] In embodiments, the operations modules 36002 include an AI system data set orchestration module 36022 that is configured and/or provided to orchestrate a set of artificial intelligence systems to perform various operations and/or tasks. For example, the AI system orchestration module 36022 may be configured and/or provided to orchestrate operation of at least a portion of the set of AI systems of the AI convergence system of systems 31000 to fulfill a set of requests. In embodiments, the AI system orchestration module 36022 interconnects two or more AI systems to

perform a particular task (*e.g.*, interconnecting the output of a first AI system, such as a pattern recognition system, to the input of a second AI system, such as a large language model that can generate summaries and/or descriptions of the detected patterns). In embodiments, the AI system orchestration module 36022 performs scheduling of the invocation of AI systems (*e.g.*, based on processing loads or queues and/or throughput of the AI systems and/or the urgency and/or priority of the requests to invoke the AI systems). In embodiments, the AI system orchestration module 36022 performs load-balancing of the requests to invoke AI systems for various tasks (*e.g.*, by monitoring a processing load and/or processing queue of each of the AI systems, and by routing requests to AI systems that have available capacity based on the load-balancing). In embodiments, the AI system orchestration module 36022 maps or remaps requested tasks to AI systems based on load-balancing, scheduling, or the like (*e.g.*, remapping a task from a first AI system that is overloaded to a second AI system that is capable of equivalently performing the task and that has available capacity). In embodiments, the AI system orchestration module 36022 adjusts allocations, provisions, and/or reservations of resources of the AI convergence system of systems to enable the AI systems to fulfill the requested invocations (*e.g.*, by acquiring more computational capacity, storage, devices such as servers or processing cores, or the like to handle an increased number of requests and/or an increased computational load of the requests, and/or by reallocating computational capacity, storage, devices such as servers or processing cores, or the like to handle a decreased number of requests and/or a decreased computational load of the requests). In embodiments, the AI system orchestration module 36022 plans and/or manages the acquisition, retention, location, installation, use, maintenance, reallocation, and/or decommissioning of the resources of the AI convergence system of systems 31000 in view of future scheduled and/or predicted demands, volumes of requests, and/or computational loads. In embodiments, the AI system orchestration module 36022 reorganizes the resources of the AI convergence system of systems 31000 based on changes in the demand for processing by the AI systems (*e.g.*, based on new types of requests that the AI systems are capable of handling, such as low-latency AI system processing requests in particular real-time contexts and/or the processing of new types of data, such as LIDAR point-cloud data). In embodiments, the AI system orchestration module 36022 reorganizes the resources of the AI convergence system of systems 31000 based on changes in the available and/or feasible AI systems (*e.g.*, allocating increased computation and/or storage for large language models and hybrid systems of increased size and/or computational complexity).

[3293] In embodiments, the AI system orchestration module 36022 may include routing data and/or interactions between and/or among data sources, clients, AI systems, hardware and/or software components, or the like. In embodiments, the AI system orchestration module 36022 determines and implements a route between a source of a request or task (*e.g.*, a client, user, requesting device, or the like) to one or more AI systems that are capable of fulfilling the request and/or performing the task (*e.g.*, AI systems that exhibit capabilities that correspond to requirements of the request and/or task). the AI system orchestration module 36022 determines and implements a route from a first AI system (*e.g.*, a generative AI system that outputs generated content) as input to a second AI system (*e.g.*, a reviewing AI system that reviews content generated

by the generative AI system). In embodiments, the AI system orchestration module 36022 determines and implements a route between a data source of data (*e.g.*, a sensor, a database, a camera, or the like) and one or more AI systems that are configured to process the data (*e.g.*, a convolutional neural network that is configured to analyze images generated by camera). In embodiments, the AI system orchestration module 36022 determines and implements a route between one or more AI systems that generate output (*e.g.*, a classifier network) and a consumer of the output (*e.g.*, a client, database, or another AI system that stores, uses, or otherwise consumes the output). In embodiments, the AI system orchestration module 36022 partitions data into portions and determines and implements routes between a source of the data and a plurality of AI systems that respectively process a portion of the data (*e.g.*, two or more AI systems that operate in parallel and/or in tandem), which may reduce latency of processing the data. In embodiments, the AI system orchestration module 36022 routes data to each of two or more AI systems, each of which may perform a same or similar type of processing (*e.g.*, redundant processing of the data by redundant AI systems) and/or different types of processing (*e.g.*, performing different tasks on the data, such as a first task of verifying the data and a second task of supplementing the data). In embodiments, the AI system orchestration module 36022 determines and implements routes between and/or among AI systems to enable interoperation to complete the task (*e.g.*, a communication path between two or more AI agents that operate independently to perform different portions of a task, such as a transaction negotiating AI agent that identifies and negotiates transactions and a transaction execution AI agent that generates and executes smart contracts to complete the transactions). In embodiments, the AI system orchestration module 36022 determines and implements routes between and/or among AI systems based on a set of available routes, such as a set of candidate communication paths by which the AI systems may communicate (*e.g.*, various network paths in a wide-area network, and/or different communication modalities such as Bluetooth, Wi-Fi, cellular, infrared, wired connections such as Ethernet, or the like). In embodiments, the AI system orchestration module 36022 determines and implements routes based on an evaluation of the requirements of a request and/or task (*e.g.*, a priority, a deadline, an amount of data, a budget, a security consideration, or the like) and corresponding features of candidate routes (*e.g.*, throughput, availability, cost, reliability, security, or the like, of various communication modalities such as Bluetooth, Wi-Fi, cellular, infrared, wired connections such as Ethernet, or the like). In embodiments, the AI system orchestration module 36022 determines and implements routes between and/or among AI systems by acquiring, purchasing, developing, reserving, provisioning, dedicating, and/or allocating one or more routes or the like (*e.g.*, reserving a communication path between two or more AI systems that are intended to communicate with high performance).

[3294] In embodiments, the operations modules 36002 include an AI system monitoring module 36024 that is configured and/or provided to analyze the performance of artificial intelligence systems while performing various operations and/or tasks. For example, the AI system monitoring module 36024 may be configured and/or provided to analyze a performance of one or more AI systems of the AI convergence system of systems 31000 to measure and compare various

performance metrics, such as accuracy, precision, recall, F1 score, bias, consistency, efficiency, latency, or the like. The AI system monitoring module 36024 may perform such evaluation using training data on which an AI system was previously trained, new training data that the AI system has not previously processed (*e.g.*, a test data set), live data that is associated with a deployment of the AI system, and/or synthetic data generated for the monitoring. The AI system monitoring module 36024 may perform such evaluation using data that is the same as or similar to data that the AI system was deployed to process and/or based on new data that the AI system was not deployed to process (*e.g.*, based on new data types that are different than those associated with the training and/or initial deployment of the AI system). The AI system monitoring module 36024 may compare the performance of an AI system at a current time with a performance of the same or similar AI system at a previous time (*e.g.*, at a time of completion of training and/or deployment). The AI system monitoring module 36024 may compare the performance of an AI system operating under a first set of conditions (*e.g.*, on a first device and/or with a first set of computational resources) with a performance of the same or similar AI system under a second set of conditions (*e.g.*, on a second device and/or with a second set of computational resources). The AI system monitoring module 36024 may compare the performance of an AI system at a current time with a simulated performance of the AI system at the current time (*e.g.*, a digital twin of the AI system). The AI system monitoring module 36024 may compare the performance of an AI system on a first data set (*e.g.*, a first set of images depicting items of a first class) with the performance of the AI system on a second data set (*e.g.*, a second set of images depicting items of a second class) The AI system monitoring module 36024 may compare the performance of an AI system with the performance of another AI system of the same or similar type (*e.g.*, comparing a performance of an artificial neural network classifier trained with a first training data set with a performance of an artificial neural network classifier trained with a second training data set). The AI system monitoring module 36024 may compare the performance of an AI system with the performance of another AI system of a different type (*e.g.*, comparing a performance of an artificial neural network classifier on a data set with a performance of a Bayesian classifier on the same data set). In embodiments, the AI system monitoring module 36024 performs comparisons using two or more performance metrics (*e.g.*, accuracy, precision, and recall) and, optionally, based on relative priorities of the performance metrics (*e.g.*, associating different relative weights and/or penalties to measurements of accuracy, precision, and recall). In embodiments, the AI system monitoring module 36024 tracks changes in the performance metrics over time and identifies trends in the performance metrics (*e.g.*, the development of AI system "drift" due to trends in the features of the data processed by the AI system, changes in the execution of the AI system, and/or the continued training of the AI system). In embodiments, the AI system monitoring module 36024 invokes a first AI system to evaluate a second AI system, and receives, from the first AI system, an indication of the performance of the second AI system. In embodiments, the AI system monitoring module 36024 monitors each AI system according to a schedule (*e.g.*, periodically testing the AI system and recording performance measurements for each period). In embodiments, the AI system monitoring module 36024 collects information from one or more other components or systems that

are associated with an AI system (*e.g.*, a sensor whose functionality depends upon the performance of an AI system) and determines a performance measurement of the AI system based on the performance of the other components or systems. In embodiments, the AI system monitoring module 36024 solicits, collects, and/or receives information from one or more entities that are associated with an AI system (*e.g.*, feedback from users or clients) and determines a performance of an AI system based on the solicited, collected, and/or received information. In embodiments, the AI system monitoring module 36024 monitors each AI system in response to a detected event or condition (*e.g.*, upon receiving an indication that a performance of an AI system may have changed, such as when an execution context of the AI system changes, or when another system detects a change in a performance of the AI system). In embodiments, the AI system monitoring module 36024 records measurements, changes, trends, and the like of the AI systems. In embodiments, the AI system monitoring module 36024 stores the measurements, changes, trends, or the like in one or more data stores or repositories, such as a database or a performance log. In embodiments, the AI system monitoring module 36024 stores the measurements, changes, trends, or the like, such as to an AI system analyzing module 36028 of the operations modules 36002.

[3295] In embodiments, the operations modules 36002 include an AI system analyzing module 36026 that analyzes the performance of the AI systems. For example, the AI system analyzing module 36026 may be configured and/or provided to analyze the performance metrics collected stored by the AI system monitoring module 36024 for various AI systems. In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system analyzing module 36026 initiates a further evaluation of the AI system to determine a cause for the change in the performance (*e.g.*, a change in the data provided as input to the AI system; a change in an execution environment of the AI system; and/or a change in a performance of the AI system due to continued training). The AI system analyzing module 36026 may use a variety of techniques to determine a cause of a change in performance, such as simulation of the AI system by one or more digital twins, Monte Carlo analysis based on sampling the effects of changes to the operation of the AI system, and/or root cause analysis that involves a determination of causal relationships among various elements of the AI system and the AI convergence system of systems 31000. In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system analyzing module 36026 changes an execution environment of the AI system to improve performance (*e.g.*, increasing computational resources allocated to the AI system, such as computational load, memory, and/or storage to improve performance). In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system analyzing module 36026 adjusts an allocation of AI systems of the AI convergence system of systems 31000 to the requests for invocations of the AI systems (*e.g.*, optionally by interoperating with the AI system orchestration module 36022) such that requests are handled by AI systems that are capable of meeting performance requirements of the requests. In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system analyzing module 36026 initiates a redevelopment of the AI system by an AI system redevelopment module 36034 (*e.g.*, a retraining and/or replacement of the AI system with another AI system).

[3296] In embodiments, the operations modules 36002 include an AI system updating module 36028 that is configured and/or provided to update deployed artificial intelligence systems. In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system updating module 36028 causes input to the AI system to be changed (*e.g.*, reformatting, annotating, curating, editing, or otherwise changing the input to the AI system to improve the performance of the AI system). In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system updating module 36028 causes an architecture of an AI system to be changed to improve the performance of the AI system (*e.g.*, increasing a context window size of an AI system to process more data, and/or increasing a number of layers of the AI system to increase learning capacity). In embodiments, in response to a detected or determined change in a performance of an AI system, the AI system updating module 36028 causes output of the AI system to be changed (*e.g.*, reformatting, annotating, curating, editing, or otherwise changing the output to the AI system to improve the performance of the AI system). In embodiments, the AI system updating module 36028 experimentally evaluates potential updates to an AI system to determine effects of the updates on the performance of the AI system (*e.g.*, experimentally increasing the computational resources allocated to AI system and measuring an effect of the increased computational resources on one or more performance metrics of the AI system). In embodiments, the AI system updating module 36028 experimentally adjusts the inputs and/or processing of an AI system, or of a simulation of the AI system (*e.g.*, of a digital twin), to determine the effect of the adjustments on the AI system. If the adjustments result in improved performance of the AI system, the AI system updating module 36028 may apply and/or commit the update to the AI system to realize and/or maintain the improvement in performance. In embodiments, in response to a detected or determined improvement in a performance of an AI system, the AI system updating module 36028 applies one or more features of the AI system that is associated with the improved performance to at least one other AI system, thereby enabling the at least one other AI system to exhibit a same or similar performance improvement.

[3297] In embodiments, the operations modules 36002 include an AI system governance module 36030 that is configured and/or provided to maintain governance of deployed artificial intelligence systems. For example, the AI system governance module 36030 may be configured and/or provided to govern one or more AI systems based on legal, social, practical, and/or regulatory policies (*e.g.*, policies relating to truthfulness and/or avoidance of harm). In embodiments, the AI system governance module 36030 monitors invocations of one or more AI systems to detect and prevent unauthorized, unintended, and/or harmful uses of the AI systems (*e.g.*, requests that are not permitted or that may contribute to harm and/or the violation of policies). In embodiments, the AI system governance module 36030 monitors input to one or more AI systems to detect and prevent unauthorized, unintended, and/or harmful inputs to the AI systems (*e.g.*, inputs that include or are associated with objectionable and/or harmful content). In embodiments, the AI system governance module 36030 monitors the processing of one or more AI systems to detect and prevent unauthorized, unintended, and/or harmful processing by the AI systems (*e.g.*, the avoidance of reasoning that may be considered deceptive or illegal). In embodiments, the AI system governance

module 36030 monitors output of one or more AI systems to detect and prevent unauthorized, unintended, and/or harmful outputs of the AI systems (*e.g.*, outputs that include or are associated with objectionable and/or harmful content). In embodiments, the AI system governance module 36030 monitors uses of the output of the AI systems by requesters and/or systems to determine unauthorized, unintended, and/or harmful uses of the AI systems (*e.g.*, malicious use of the processing capabilities and/or output of the AI systems). In embodiments, the AI system governance module 36030 responds to detected requests, inputs, processing, outputs, and/or uses of the AI systems in various ways, such as modifying, censoring, editing, redacting, and/or blocking a request, input, processing, output, and/or use of the AI system. In embodiments, the AI system governance module 36030 receives, generates, adjusts, and/or reports policies based on the monitoring and governance, such as generating guidelines and/or tests to evaluate the requests, inputs, processing, outputs, and/or uses of the AI systems that may be associated with malicious, unauthorized, unintended, and/or harmful content or activities. In embodiments, the AI system governance module 36030 generates, adapts, and/or monitors one or more governance mechanisms applied to an AI system, such as monitoring an effectiveness of a rule set applied to the requests and/or inputs of an AI system to ensure that the processing and/or outputs of the AI system do not include unauthorized, unintended, and/or harmful content. In embodiments, the AI system governance module 36030 is configured to detect attempts to circumvent, subvert, remove, or otherwise avoid governance mechanisms applied to one or more AI systems (*e.g.*, attempts to present input that is unauthorized, unintended, and/or harmful, but that is disguised to avoid detection by a governance rule set) and to update the policies and/or governance mechanisms of the AI systems to detect, reduce, alleviate, prevent, report, and/or penalize such avoidance. In embodiments, the AI system governance module 36030 is configured to detect attempts to violate the security, integrity, and/or availability of the AI systems (*e.g.*, denial-of-service attacks, AI system "poisoning" attacks, and/or AI system exploration, eavesdropping, or data exfiltration) and to take steps to reduce, alleviate, prevent, report, and/or penalize such attacks.

[3298] In embodiments, the AI system governance module 36030 configures and/or provides supervision of one or more AI systems. For example, the AI systems may include one or more generative AI systems, and the AI system governance module 36030 may supervise the generation of content to promote positive features (*e.g.*, quality, accuracy, coherence, or the like) and/or to reduce or prevent negative features (*e.g.*, inaccuracy, incoherence, harmfulness, or the like). The AI system governance module 36030 may apply such supervision to the inputs to the generative AI systems (*e.g.*, prompts and/or content sources provided as input), to the processing of the generative AI systems (*e.g.*, to particular forms of information and/or reasoning used by the generative AI system), to the output of the generative AI systems (*e.g.*, to content generated based on prompts and/or input), and/or to the uses of the generative AI systems (*e.g.*, to the purposes for which generated content is used). The AI system governance module 36030 may apply such supervision based on rules, guidelines, heuristics, principles, comparisons with positive and negative content and/or processing, or the like. The AI system governance module 36030 may apply such supervision by generating and/or applying a supervising AI system to another AI system

(*e.g.*, a classifier network that classifies the output of a generative AI system as potentially negative). AI system governance module 36030 may apply such supervision with input from human reviewers (*e.g.*, the AI system governance module 36030 may involve humans for spot-checking content, for scoring or rating content, for reviewing content generated by the AI systems that has been flagged as potentially negative, and/or for responding to reports of negative content generated by the AI systems). The AI system governance module 36030 may use the results of automated and/or human supervision to modify generated content; to modify the AI systems (*e.g.*, training, retraining, and/or replacing the AI systems); to create, modify, and/or refine the supervision (*e.g.*, modifying the rules, heuristics, principles, or the like to prevent the generation of further negative content by the AI systems); to create, modify, and/or refine supervising AI systems that are connected to other AI systems (*e.g.*, training, retraining, and/or replacing a supervising AI system with an improved supervising AI system); and/or to generate reports about the AI systems (*e.g.*, reports of the prompts, performance, types, and/or uses of content generated by the AI systems).

[3299] In embodiments, the operations modules 36002 include an AI system redevelopment module 36032 that is configured and/or provided to redevelop deployed artificial intelligence systems. For example, in response to the AI system analyzing module 36026 detecting reduced performance of an AI system, the AI system redevelopment module 36032 may be configured and/or provided to initiate retraining of the AI system using additional training data (*e.g.*, additional synthetic training data generated by the AI system training data set augmentation module 36008 corresponding to new data that the AI system has not previously been trained to receive and process). In embodiments, in response to the AI system analyzing module 36026 detecting reduced performance of an AI system, the AI system redevelopment module 36032 causes the AI system training module 36010 to train a replacement AI system, and causes the AI system verifying module 36012 to verify that the replacement AI system addresses the reduced performance of the AI system. Based on such verifying, the AI system redevelopment module 36032 may cause the AI system deployment module 36018 to deploy the replacement AI system to substitute for the original AI system. In embodiments, the AI system redevelopment module 36032 may experimentally task the AI system training module 36010 with developing candidate AI systems that might replace an existing AI system. As compared with an existing AI system, a candidate AI system may be trained on different training data, with a different training regimen, with a different AI system type and/or architecture. A candidate AI system may be trained in a same or similar manner as an existing AI system, but may result in different and possibly improved performance due to the stochastic nature of AI system training. The AI system redevelopment module 36032 may initiate the redevelopment and/or experimental testing of candidate AI systems in response to a change of priority of tasks related to an AI system (*e.g.*, an updated indication that a particular type of error is more harmful and/or more common than was previously initiated). The AI system redevelopment module 36032 may initiate the redevelopment and/or experimental testing of candidate AI systems in response to a change in the resources of the AI convergence system of systems 31000 (*e.g.*, increased computational capacity may enable a candidate AI system with

greater learning capacity to outperform an existing AI system). The AI system redevelopment module 36032 may initiate the redevelopment and/or experimental testing of candidate AI systems in response to advances in AI systems research or technology (*e.g.*, a development of new model types, a new training regimen, new types of AI system hardware and/or software, or scientific discoveries relating to AI system performance and/or performance optimization techniques). Based on a candidate AI system exhibiting improved performance relative to an existing AI system, the AI system redevelopment module 36032 may initiate a deployment of the candidate AI system to replace the existing AI system.

[3300] In embodiments, the operations modules 36002 include an AI system logging module 36034 that is configured and/or provided to log the activities of artificial intelligence systems and/or operations modules 36002 while performing various operations and/or tasks. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system generation module 36004 in generating and/or storing new AI systems. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system training data set generation module 36006 in generating training data sets for the training of various AI systems. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the training data set augmentation module 36008 in the number, kinds, sources, techniques, and examples of synthetic data generated for various training data sets, as well as the circumstances for such augmentation (*e.g.*, the issues with a training data set that synthetic data was generated to address). In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system training module 36010 in selecting, initiating, monitoring, validating, and completing the training of the AI systems of the AI convergence system of systems 31000, as well as the resulting performance of the trained AI systems (*e.g.*, the performance, accuracy, recall, F1 score, bias, consistency, efficiency, and/or latency of respective trained AI systems). In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system verifying module 36012 in verifying the generated AI systems (*e.g.*, measurements of the performance of trained AI models on test data sets). In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system adaptation module 36014 in adapting AI systems for particular devices, execution environments, data types, tasks, contexts, use cases, or the like. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system aggregation module 36016 in aggregating AI systems to generate and interconnect two or more AI systems to form hybrids, ensembles, sequences or chains, managed AI systems, AI system swarms, or the like. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system deployment module 36018 to deploy AI systems in response to various requests, which may originate within the AI convergence system of systems 31000 or external to the AI convergence system of systems 31000. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system invocation module 36020 to invoke one or more AI systems to

fulfill one or more requests. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system orchestration module 36022 to organize the invocations of the AI systems of the AI convergence system of systems 31000 to meet the requirements of the requests and/or tasks, such as changes due to load-balancing or scheduling. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system monitoring module 36024 to monitor the performance of the AI systems of the AI convergence system of systems 31000, such as the planned and/or completed schedule of measuring and monitoring the performance of the AI systems and/or the results of the performance measurements. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system analyzing module 36026 to analyze the measured performance of the AI systems and the steps taken by the AI system analyzing module 36026 or other modules to respond to changes in performance (*e.g.*, retraining AI systems, replacing AI systems, and/or adapting some AI systems to include features that resulted in performance improvements of other AI systems). In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system updating module 36028 to update various AI systems in response to the analyses performed by the AI system analyzing module 36028. In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system governance module 36030 to manage and/or govern the AI systems based on legal, social, practical, and/or regulatory policies (*e.g.*, policies relating to truthfulness and/or avoidance of harm). In embodiments, the AI system logging module 36034 is configured and/or provided to detect and log the operations of the AI system redevelopment module 36032 to redevelop the AI systems based on the performance of the AI systems, the resources of the AI convergence system of systems 31000, and/or the development of new AI systems. In embodiments, the AI system logging module 36034 correlates the activities of various modules to determine causal relationships between the operations modules 36002 (*e.g.*, a first event involving a deployment of an AI system in an environment that is associated with a second event involving a performance measurement of another AI system in the same environment; a first event involving an adaptation of an AI system followed by a second event involving a performance measurement of the adapted AI system; or a first event involving a reorganization of the AI convergence system of systems 31000 by the AI system orchestration module 36022 followed by a second event involving a change in a performance measurement of an AI system by the AI system monitoring module 36024). In embodiments, the AI system logging module 36034 logs the operations of the AI models and/or the operations modules 36002 in one or more event logs.

Operations Layer – Example Use Cases

[3301] In embodiments, the operations modules 36002 are configured to serve one or more smart asset operations. For example, an asset may comprise a fund, an amount of currency such as an accumulation of cryptocurrency, a digital contract, a valuable object, or the like. Such assets may be made "smart" by supplementing the asset with intelligent functions that resemble cognition and that improve the features of the asset, such as value, security, and/or longevity. For example, a

fund may be made "smart" by predicting or detecting events that may positively or negatively affect the value of the fund, identifying steps that may be taken responsive to the events, and/or executing the steps. A digital contract may be made "smart" by predicting or detecting events that may positively or negatively affect the execution of the smart contract by one or more parties, identifying steps that may be taken responsive to the events, and/or executing the steps. A valuable object may be made "smart" by predicting or detecting events that may positively or negatively affect the value of the object, identifying steps that may be taken responsive to the events, and/or executing the steps. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features for such assets. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to an asset, such as an AI system that monitors social, political, economic, scientific, and/or technical events and extrapolates positive and/or negative effects of the current events on a current and/or future value of the asset. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to an asset, such as an AI system that identifies actions that can be taken regarding the asset in response to social, political, economic, scientific, and/or technical events. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to an asset, such as an AI system that executes actions to promote the value, security, and/or longevity of the asset in response to social, political, economic, scientific, and/or technical events.

[3302] In embodiments, the operations modules 36002 are configured to serve one or more deployments of smart infrastructure. For example, the infrastructure of an organization such as a government, a university, a hospital, or the like may include a set of interconnected infrastructure systems, such as power, communication, information technology, inventory management, transportation, maintenance, cleaning, and auditing. Each infrastructure system may include a set of resources, such as power infrastructure including power storage, wiring, regulation such as circuit breakers, switches for failover and load-balancing, and the like. The infrastructure may be made "smart" by supplementing the infrastructure with intelligent functions that resemble cognition and that improve the functionality of the infrastructure, such as availability, security, efficiency, and/or maintenance. For example, power infrastructure may be made "smart" by predicting and proactively addressing potential power failures due to events such as weather patterns, climate change, events in energy markets, and/or changes in the availability of energy-related supplies such as wiring and batteries. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to infrastructure. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to infrastructure, such as an AI system that monitors social, political, economic, scientific, and/or technical events and extrapolates positive and/or negative effects of the current events on a current and/or future operation of various systems of the infrastructure. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the infrastructure, such as an

AI system that identifies actions that can be taken to promote and/or maintain the availability, security, efficiency, and/or maintenance of the infrastructure system. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the infrastructure, such as an AI system that executes actions to promote the availability, security, efficiency, and/or longevity of the infrastructure in response to social, political, economic, scientific, and/or technical events.

[3303] In embodiments, the operations modules 36002 are configured to serve one or more smart machines. For example, an industrial machine in a mining facility may include a number of mining-related functions, such as extracting a mineral, refining a mineral, and assessing the refined mineral for quality, quantity, purity, and/or value. The machine may be made "smart" by supplementing the machine with intelligent functions that resemble cognition and that improve the functionality of the machine, such as performance, efficiency, and/or longevity. For example, the machine may be made "smart" by adapting the manner of extracting minerals to promote the speed, efficiency, completeness, security, and/or integrity of the extraction process. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the machine. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the machine, such as an AI system that predicts how changes in the mine and/or other machines of the mine may impact the current and/or future operation of the machine. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the machine, such as an AI system that identifies actions that can be taken to promote and/or maintain the performance, efficiency, and/or longevity of the machine (*e.g.*, physical changes or upgrades of the machine, changes to the manner and/or timing of operation of the machine, and/or maintenance actions regarding various components of the machine). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the machine, such as an AI system that executes actions to promote the performance, efficiency, and/or longevity of the machine (*e.g.*, automatically arranging maintenance of components of the machine that are at risk of wear and/or failure).

[3304] In embodiments, the operations modules 36002 are configured to serve one or more robotic operations. For example, a robotic operation may include a fleet of robots, each having a set of capabilities (*e.g.*, sensors, modes of transportation, tools, end effectors, and the like) and requirements (*e.g.*, fuel or power storage, protection from elements such as water or electrical shock, and periodic maintenance), and each being assigned to perform one or more tasks (*e.g.*, transporting materials, monitoring a location, and/or serving a user). The robotic operation may be made "smart" by supplementing the robotic operation with intelligent functions that resemble cognition and that improve the functionality of the robotic operation, such as the assignment of robots to tasks, the maintenance and upgrading of the robotic fleet, and the performance of tasks by the robots. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the robotic operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI

systems to add a predictive "smart" feature to the robotic operation, such as an AI system that predicts how changes in the location, orientation, environment, climate, cargo, and/or containers of the robotic operation may impact the performance of the tasks by the robotic fleet, such as obstacles or interruptions of service that might prevent the robots from completing one or more tasks. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the robotic operation, such as an AI system that identifies actions that can be taken to promote the performance, reliability, efficiency, and/or completion of the tasks by the robotic fleet, including adapting the assignment, scheduling, and/or prioritization of the robots of the robotic fleet to the set of available tasks. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the robotic operation, such as an AI system that executes actions to promote the performance, reliability, efficiency, and/or completion of the tasks by the robotic fleet, such as automatically deploying, assigning, reassigning decommissioning, and/or upgrading robots to perform the tasks of the robotic operation, and/or automatically maintaining and/or upgrading the robotic fleet to preserve and/or add capabilities to the robotic fleet.

[3305] In embodiments, the operations modules 36002 are configured to serve one or more autonomous vehicle operations. For example, an autonomous vehicle operation may include a number of autonomous vehicles, such as taxis, buses, ambulances, trucks, trains, monorails, trolleys, aircraft, ferries, and cargo ships, as well as resources such as garages, fuel depots, roads, waterways, railways, and airspaces. The autonomous vehicle operation may coordinate the autonomous vehicles and resources to arrange the pickup, transport, and delivery of drivers, operators, passengers, clients, and/or cargo, and may provide a number of features for the drivers, operators, passengers, clients, and/or cargo, such as physical isolation and protection from weather elements, comfort and privacy of the drivers, operators, passengers, clients, and/or cargo, and physical security of the drivers, operators, passengers, clients, and cargo during transit. The autonomous vehicle operation may be made "smart" by supplementing the autonomous vehicle operation with intelligent functions that resemble cognition and that improve the functionality of the autonomous vehicle operation, such as the selection of transportation routes that improve the speed, safety, and efficiency of transporting the drivers, operators, passengers, clients, and/or cargo. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the autonomous vehicle operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the autonomous vehicle operation, such as an AI system that predicts how changes in the environment, population, and/or maintenance of transit resources (*e.g.*, roads or railways), autonomous vehicles, drivers, operators, passengers, clients, and/or cargo may impact the safe and/or timely transportation of the drivers, operators, passengers, clients, and/or cargo, such as an arrival time at a destination and/or a condition of the drivers, operators, passengers, clients, and/or cargo upon arrival. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the autonomous vehicle operation, such as an AI system that identifies actions that can be taken

to promote and/or maintain the protection, security, and/or maintenance of the autonomous vehicles, drivers, operators, passengers, clients, and/or cargo (*e.g.*, scheduling a stop along a route for the rest, relaxation, maintenance, and/or fulfillment of needs of the drivers, operators, passengers, clients, and/or cargo). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the autonomous vehicle operation, such as an AI system that executes actions to protection, security, and/or maintenance of the autonomous vehicles, drivers, operators, passengers, clients, and/or cargo (*e.g.*, automatically updating a route of an autonomous vehicle to improve an estimated arrival time at a destination and/or to promote the rest, relaxation, maintenance, and/or fulfillment of needs of the drivers, operators, passengers, clients, and/or cargo).

[3306] In embodiments, the operations modules 36002 are configured to serve one or more UAV operations. For example, a UAV operation may include a fleet of unmanned aerial vehicles, such as drones, light aircraft, helicopters, balloons, dirigibles, airships, rockets, satellites, or the like. The unmanned aerial vehicles may be autonomously controlled by onboard or remote AI systems to perform the takeoff, route selection and correction, airborne maneuvers, status monitoring, and landing of the aircraft. The UAV operation may assign various UAVs of the UAV fleet to various tasks, such as monitoring or measuring a particular geographic region or a subject such as a field of crops, livestock, or an industrial or military operation; deploying substances such as water, fire suppression substances, fertilizer, or pesticide; transporting cargo between various locations; delivering fuel and/or power to other UAVs; or providing first-responder services such as transporting first responders or equipment to address health emergencies or public safety events. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the UAV operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the UAV operation, such as an AI system that predicts how changes in the environment, climate, traffic, wildlife, population, and/or equipment may impact the performance of the tasks by the UAV fleet, such as obstacles or interruptions of service that might prevent the UAVs from completing one or more tasks. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the UAV operation, such as an AI system that identifies actions that can be taken to promote the performance, reliability, efficiency, and/or completion of the tasks by the UAV fleet, including adapting the assignment, scheduling, and/or prioritization of the UAVs of the UAV fleet to the set of available tasks. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the UAV operation, such as an AI system that executes actions to promote the performance, reliability, efficiency, and/or completion of the tasks by the UAV fleet, such as automatically deploying, assigning, reassigning decommissioning, and/or upgrading UAVs to perform the tasks of the UAV operation, and/or automatically maintaining and/or upgrading the UAV fleet to preserve and/or add capabilities to the UAV fleet.

[3307] In embodiments, the operations modules 36002 are configured to serve one or more additive manufacturing operations. For example, an additive manufacturing operation may include

an array of 3D printers, each configured to receive a three-dimensional model and to generate one or more physical objects corresponding to the three-dimensional model via additive manufacturing (*e.g.*, fused deposition modeling (FDM) and/or resin printing). Each 3D printer may include capabilities (*e.g.*, print speeds, print quality, temperatures ranges, material types, and object sizes) and may provide various functions during printing (*e.g.*, monitoring, troubleshooting, adaptation of printing to accommodate various materials and/or techniques, object post-processing, object assembly, and/or object delivery). The additive manufacturing operation may receive a set of work orders for various objects and may assign each work order to one or more 3D printers by matching the requirements of each work order to the capability and/or availability of the respective 3D printers. The additive manufacturing operation may be made "smart" by supplementing the additive manufacturing operation with intelligent functions that resemble cognition and that improve the functionality of the additive manufacturing operation, such as the assignment of 3D printers to work orders, the maintenance and upgrading of the 3D printers, and the execution of work orders by the 3D printers. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the additive manufacturing operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the additive manufacturing operation, such as an AI system that predicts how changes in the environment or climate, the availability and condition of 3D printing materials, and the conditions of the 3D printers may impact the execution of the work orders by the 3D printers, such as obstacles or interruptions of service that might prevent the 3D printers from completing one or more work orders. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the additive manufacturing operation, such as an AI system that identifies actions that can be taken to promote the performance, reliability, efficiency, and/or completion of the work orders by the 3D printers, including adapting the assignment, scheduling, and/or prioritization of the work orders by the 3D printers. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the additive manufacturing operation, such as an AI system that executes actions to promote the performance, reliability, efficiency, and/or completion of the work orders by the 3D printers, such as automatically deploying, assigning, reassigning decommissioning, and/or upgrading 3D printers to perform the tasks of the additive manufacturing operation, and/or automatically maintaining and/or upgrading the 3D printers fleet to preserve and/or add capabilities to the additive manufacturing operation.

[3308] In embodiments, the operations modules 36002 are configured to serve one or more human-machine cooperative operations. For example, a human-machine cooperative operation may involve one or more tasks that are performed by one or more humans cooperating with one or more machines in fields such as mining, manufacturing, surveying, construction, healthcare, education, emergency response, or the like. A human and a machine may cooperate to perform a physical task together, such as carrying a piece of equipment, furniture, or a material such as a pane of glass. A human and a machine may cooperate to choose a plan for completing a task, such

as surveying a geographic region to measure a population of wildlife. A human and a machine may cooperate by the human supervising the performance of a task by a machine, such as a robot engaging in a manufacturing or construction operation while a human monitors the operation of the robot to detect and address problems. A human and a machine may cooperate by the robot observing the human during the performance of a task to learn how to perform the task, such as a robotic process automation (RPA) task that enables the robot to perform the task autonomously in the future. A human and a machine may cooperate by the human observing the robot during the performance of a task to learn how to perform the task, such as a skill-sharing operation where a robot teaches a human to perform a task, which may have been learned from another human. The human-machine cooperative operation may be made "smart" by supplementing the human-machine cooperative operation with intelligent functions that resemble cognition and that improve the functionality of the human-machine cooperative operation, such as the mapping of tasks to combinations of one or more humans and one or more machines. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the human-machine cooperative operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the human-machine cooperative operation, such as an AI system that predicts which humans and which machines will cooperate well together to perform a particular task. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the human-machine cooperative operation, such as an AI system that identifies the occurrence and/or causes of current and/or future disagreements between a human cooperating with a machine. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the human-machine cooperative operation, such as an AI system that executes actions to address disagreements between a human cooperating with a machine, such as conversations, mediation by the AI system or by another human and/or machine, clarification of communication between the human and the machine, review and analysis of previous interaction between the human and the machine that resulted in the disagreement, and/or adjustments of the assignment of humans and machines to various cooperative tasks.

[3309] In embodiments, the operations modules 36002 are configured to serve one or more industrial operations. For example, an industrial operation may involve one or more mining, manufacturing, construction, and/or chemical and/or biologic synthesis operations that result in products such as consumer goods, food, clothing, electronics, furniture, vehicles, household chemicals, pharmaceuticals, biologic strains, or the like. The industrial process may involve various industrial tasks such as receiving, storing, processing, manipulating, assembling, packaging and/or maintaining various items, compositions, biologic strains, or the like. The industrial operation may be made "smart" by supplementing the industrial operation with intelligent functions that resemble cognition and that improve the functionality of the industrial operation, such as the performance of the tasks involved in the industrial operation. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or

invoke one or more AI systems to provide "smart" features to the industrial operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the industrial operation, such as an AI system that predicts changes in the supply chain that may positively or negatively impact the availability, quality, and/or suitability of raw materials for the tasks of the industrial operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the industrial operation, such as an AI system that identifies actions that can be taken to avoid complications arising from events, such as stockpiling an excess of certain raw materials involved in the industrial operation to address a progressing and/or predicted supply chain issue. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the industrial operation, such as an AI system that executes transactions to acquire an excess of certain raw materials involved in the industrial operation to address a progressing and/or predicted supply chain issue.

[3310] In embodiments, the operations modules 36002 are configured to serve one or more automated machine monitoring operations. For example, an automated machine monitoring operation may involve monitoring operations of a machine to detect the occurrence, progress, and/or quality of the performance of tasks, the status of the machine (*e.g.*, temperature, inventory of consumables, and/or freedom from defects and/or anomalous behavior), and interactions between the machine and other machines of a facility. The automated machine monitoring operation may be made "smart" by supplementing the automated machine monitoring operation with intelligent functions that resemble cognition and that improve the functionality of the automated machine monitoring operation, such as monitoring the completion of tasks by the machine, monitoring the status of the machine, and monitoring the interactions between the machine and other machines of a facility. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the automated machine monitoring operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the automated machine monitoring operation, such as an AI system that predicts how changes to a facility (*e.g.*, the locations, orientations, and byproducts of other machines of a facility) may positively or negatively impact the monitoring of the performance, status, or interaction of a machine with other machines of a facility (*e.g.*, organization of a facility that resulted in a blockage of a view of a machine by a monitoring camera and/or interference with the monitoring of the machine by a sensor). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the automated machine monitoring operation, such as an AI system that identifies causes of problems with the monitoring of the performance of tasks by the machine, the status of the machine, and/or the interaction of the machine with other machines of the facility (*e.g.*, determining that the location of materials in a facility interferes with a monitoring of a machine of the facility by the AI system). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the automated machine monitoring operation, such as an AI system that executes

actions to address, reduce, and/or prevent problems with the monitoring of the performance of tasks by the machine, the status of the machine, and/or the interaction of the machine with other machines of the facility (*e.g.*, autonomously instructing a robot of the facility to relocate materials in the facility to address a blockage of a monitoring process).

[3311] In embodiments, the operations modules 36002 are configured to serve one or more maintenance operations. For example, a maintenance operation may involve surveying, testing, resupplying, repairing, reprogramming, refreshing, cleaning, and/or replacing a machine or one or more components of a machine. The maintenance operation may be made "smart" by supplementing the maintenance operation with intelligent functions that resemble cognition and that improve the functionality of the maintenance operation, such as refining the maintenance operations for improved efficiency, speed, availability, reliability, completion, and/or recordation of the maintenance operations. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the maintenance operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the maintenance operation, such as an AI system that predicts how changes to the machine or an environment of the machine may positively or negatively impact the relevance, urgency, timeliness, extent, and/or completeness of maintenance operations for the machine.

[3312] In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the maintenance operation, such as an AI system that improves the diagnostic accuracy, reliability, and/or efficiency of the maintenance operations of the machine in view of changes to the machine, uses of the machine (*e.g.*, the materials and/or tasks for which the machine is used), and/or an environment of the machine. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the maintenance operation, such as an AI system that executes actions to improve the diagnostic accuracy, reliability, and/or efficiency of the maintenance operations, such as automatically adjusting a maintenance schedule in response to changes to the machine, uses of the machine (*e.g.*, the materials and/or tasks for which the machine is used), and/or an environment of the machine.

[3313] In embodiments, the operations modules 36002 are configured to serve one or more robotic process automation (RPA) operations. For example, an RPA operation may involve a learning software process that observes a human during the performance of a task (*e.g.*, a physical task, a cognitive process, and/or an interaction between the human and one or more other humans, robots, and/or software processes such as user interfaces) in order to learn to perform the task in a corresponding and/or equivalent autonomous manner using software, hardware, and/or robotic processes. The RPA operation may be made "smart" by supplementing the RPA operation with intelligent functions that resemble cognition and that improve the functionality of the RPA operation, such as improving the accuracy, clarity, reliability, and/or situational awareness of the learning software process while learning and/or autonomously performing the task. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or

invoke one or more AI systems to provide "smart" features to the RPA operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the RPA operation, such as an AI system that predicts deficiencies in the learning process, such as features of the task that may not be or may not have been adequately conveyed to and/or learned by a learning software process (*e.g.*, lack of clarity, incomplete explanations, confusion, inconsistency, and/or incomplete records of the performance of the task by the human). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the RPA operation, such as an AI system that identifies ways to address deficiencies in the learning process (*e.g.*, requests that may be presented to the human, such as requests to explain certain aspects of the task or the context of the task and/or requests for an additional performance of the task). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the RPA operation, such as an AI system that executes actions to address, reduce, and/or prevent problems with the performance of tasks by the machine, the status of the machine, and/or the interaction of the machine with other machines of the facility.

[3314] In embodiments, the operations modules 36002 are configured to serve one or more converged information technology / operational technology (IT/OT) operations. For example, a converged IT/OT operation may involve a combination of a traditional, general-purpose information technology environment and resource set (*e.g.*, general-purpose networking, data storage, software and device deployment, and an information security model) with operation-specific resources and operations (*e.g.*, specific industrial machines, controllers, communication protocols, and monitoring). The converged IT/OT operation may be made "smart" by supplementing the converged IT/OT operation with intelligent functions that resemble cognition and that improve the functionality of the converged IT/OT operation, such as improving the synergy, adaptability, reliability, and robustness of the mapping of IT resources to specific industrial operations and resources. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the converged IT/OT operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the converged IT/OT operation, such as an AI system that predicts how changes to the industrial environment may affect the industrial operations and the assignment of information technology resources thereto (*e.g.*, increases of new industrial machines and/or workloads that could potentially create shortages of information technology resources, such as network capacity and/or storage capacity). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the converged IT/OT operation, such as an AI system that identifies opportunities to acquire, map, provision, reserve, deploy, configure, and/or use generic IT resources for specific industrial operations in order to provide capabilities, improve features, and/or address potential problems of the industrial operations such as yield, quality, and/or efficiency. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the converged IT/OT operation, such as an AI

system that executes actions to address, reduce, and/or prevent problems with the mapping of information technology resources to specific industrial operations, such as automatically acquiring and/or reassigning information technology resources to the industrial operations based on demand, requirements, priorities, timing, side-effects, statuses, and/or outcomes of the industrial operations.

[3315] In embodiments, the operations modules 36002 are configured to serve one or more insurance, underwriting, and/or risk management operations. For example, an insurance, underwriting, and/or risk management operation may involve a determination of one or more risks related to an entity, resource, asset, operation, venture, objective, or the like, wherein the risks may affect a condition, value, availability, suitability, longevity, functionality, and/or reputation of the entity, resource, asset, operation, venture, objective, or the like. The insurance, underwriting, and/or risk management operation may also involve an assessment of an overall risk profile, options that may mitigate such risks, and/or the availability, security, reliability, cost, and/or value of insurance and/or underwriting of such entity, resource, asset, operation, venture, objective, or the like, including options for predicting, detecting, mitigating, compensating, and/or alleviating such risks. The insurance, underwriting, and/or risk management operation may be made "smart" by supplementing the insurance, underwriting, and/or risk management operation with intelligent functions that resemble cognition and that improve the functionality of the insurance, underwriting, and/or risk management operation, such as improving the accuracy, likelihood, clarity, extent, and/or outcomes of predicted risks and/or the availability, security, reliability, cost, and/or value of insurance and/or underwriting of such risks. In such contexts, and in embodiments, the operations modules 36002 are configured to generate and/or invoke one or more AI systems to provide "smart" features to the insurance, underwriting, and/or risk management operation. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a predictive "smart" feature to the insurance, underwriting, and/or risk management operation, such as an AI system that predicts how social, political, economic, scientific, and/or technical developments are likely to affect predictions of risks and/or insurance and/or underwriting thereof. In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add a reactive "smart" feature to the insurance, underwriting, and/or risk management operation, such as an AI system that recommends adjustments to an entity, resource, asset, operation, venture, objective, or the like to reduce or mitigate a risk profile in view of social, political, economic, scientific, and/or technical developments, and/or adjustments to an insurance and/or underwriting policy (*e.g.*, informing individuals of changes to risk profiles regarding an entity, resource, asset, operation, venture, objective, or the like, and/or generating and presenting recommendations to adjust a scope, redundancy, and/or cost of coverage by insurance and/or underwriting in response to social, political, economic, scientific, and/or technical developments). In embodiments, the operations modules 36002 generate and/or invoke one or more AI systems to add an executive "smart" feature to the insurance, underwriting, and/or risk management operation, such as an AI system that executes actions to adjust an entity, resource, asset, operation, venture, objective, or the like (*e.g.*, changing a level of security of an entity, resource, asset, operation, venture, objective, or the like, and/or automatically executing transactions to scale a level, scope, redundancy, and/or

cost of coverage by insurance and/or underwriting in response to social, political, economic, scientific, and/or technical developments).

Network Layer

Network modules

[3316] Referring to FIG. 271 and FIG. 275, the AI convergence system of systems 31000 may include a network layer 37000 having a set of network modules 37002. In embodiments, the network modules 37002 may include an adaptive network routing module 37004, an adaptive error correction module 37006, an adaptive network storage module 37008, an adaptive protocol selection module 37010, an adaptive block size module 37012, an adaptive transmission window module 37014, an edge networking module 37016, a cloud networking module 37018, a communication network 37020, a cellular module 37022, a WiFi module 37024, an ORAN module 37026, a Bluetooth module 37028, a satellite module 37030, an IoT network module 37032, a mesh network module 37034, a social network module 37036, and an insurance network module 37038. In embodiments, the set of network modules 37002 may further include a power grid module, a shipping network module, a delivery network module, a distributed energy network module, a healthcare network module, an IoT-in-a-box module, an industrial IT/OT module, a transaction network module, a financial network module, a payments network module, a transportation network module, a value chain network module, and a supply chain network module.

[3317] In some embodiments, each network module of the network layer 37000 may include one or more processors and one or more memories storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to perform an intelligence task related to that particular network module. Such intelligence tasks may relate to optimization tasks, generative tasks, predictive tasks, detection and/or identification tasks, decision support tasks, automation tasks, configuration and/or control tasks, and the like.

[3318] The machine learning systems may be trained using supervised, unsupervised, or reinforcement learning techniques on diverse types of data, including simulated data (e.g., data generated from a simulation engine, a foundation world model, a set of digital twins, or the like). The neural networks may comprise multiple layers including, but not limited to, convolutional layers, recurrent layers, and fully connected layers. The artificial intelligence systems may implement various algorithms including decision trees, random forests, and gradient boosting machines. The network layer 37000 may further include a simulation engine or set of foundation world models operating in parallel with the machine learning systems, artificial intelligence systems, and/or neural networks, wherein the simulation engine and/or set of foundation world models are configured to generate real-time simulations of scenarios associated with each intelligence task, validate decisions, and provide feedback to the machine learning models for continuous improvement and adaptation to changing conditions. In embodiments, the simulation engine may employ physics-based modeling, discrete event simulation, or agent-based modeling techniques to accurately represent the dynamics of real-world behavior.

Adaptive network routing module

[3319] For example, an adaptive network routing module 37004 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize network routing paths in real time. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting traffic patterns across nodes, detection and identification tasks to recognize network bottlenecks or anomalies, automation tasks to dynamically adjust routing configurations based on current network conditions, decision support tasks by recommending optimal paths for high-priority data and configuration tasks to enforce changes across the network infrastructure to ensure efficient data flow and minimized latency, and the like.

Adaptive error correction module

[3320] In another example, an adaptive error correction module 37006 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to identify, predict, and/or correct errors in network communication in real time. Additional and/or alternative intelligence tasks may involve predictive tasks like anticipating potential error-prone scenarios based on historical data, detection and identification tasks to isolate and classify errors in data packets or transmissions, automation tasks to dynamically apply error correction protocols or retransmission strategies, decision support tasks to recommend optimal recovery mechanisms for maintaining data integrity, and configuration tasks to adjust system parameters to prevent future errors, among many others.

Adaptive network storage module

[3321] In yet another example, an adaptive network storage module 37008 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize the allocation, retrieval, and management of storage resources across the network. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting storage capacity requirements based on historical usage patterns, detection and identification tasks to monitor and address storage bottlenecks or underutilized resources, automation tasks to dynamically allocate or reallocate storage to maintain optimal performance, decision support tasks to recommend data replication or tiering strategies for improved efficiency and reliability, configuration tasks to enforce policies for storage redundancy and access control, and the like.

Adaptive protocol selection module

[3322] In yet another example, an adaptive protocol selection module 37010 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to dynamically select the most appropriate communication protocols based on real-time network conditions and requirements. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting protocol performance under varying loads, detection and identification tasks to assess protocol compatibility or inefficiencies, automation tasks to switch protocols during transmission, decision

support tasks to recommend protocols for specific scenarios, configuration tasks to optimize protocol settings for maximum efficiency, and the like.

Adaptive block size module

[3323] In yet another example, an adaptive block size module 37012 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to determine and adjust block sizes dynamically to optimize data transfer efficiency. Additional and/or alternative intelligence tasks may involve predictive tasks like estimating optimal block sizes based on network latency and bandwidth, detection and identification tasks to flag inefficiencies caused by suboptimal block sizes, automation tasks to adjust block sizes in real time during data transfers, decision support tasks to recommend configurations for specific network conditions, configuration tasks to enforce block size policies across the network, and the like.

Adaptive transmission window module

[3324] In yet another example, an adaptive transmission window module 37014 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to dynamically adjust transmission windows to optimize data flow and minimize congestion. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting optimal transmission windows based on traffic patterns, detection and identification tasks to identify bottlenecks or excessive retransmissions, automation tasks to modify transmission windows in real time to balance throughput and reliability, decision support tasks to suggest window sizes for specific workloads, configuration tasks to apply policies for transmission efficiency, and the like.

Edge networking module

[3325] In yet another example, an edge networking module 37016 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to enable localized data processing and communication at the network edge. Additional and/or alternative intelligence tasks may involve predictive tasks like anticipating traffic loads at edge nodes, detection and identification tasks to monitor edge device performance or connectivity issues, automation tasks to redistribute workloads between edge and central resources, decision support tasks to recommend edge deployment strategies, configuration tasks to manage edge resource allocation, and the like.

Cloud networking module

[3326] In yet another example, a cloud networking module 37018 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize the connectivity and resource utilization of cloud-based infrastructures. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting bandwidth demands for cloud applications, detection and identification tasks to diagnose cloud network performance issues, automation tasks to balance traffic between virtual network nodes, decision support tasks to optimize network configurations for cloud workloads, configuration tasks to manage cloud network policies, and the like.

Communication network

[3327] In yet another example, a communication network module 37020 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to manage and optimize end-to-end data communication across the network. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting communication delays or outages, detection and identification tasks to locate and resolve data packet loss, automation tasks to reroute traffic for maximum efficiency, decision support tasks to enhance network configurations for specific use cases, configuration tasks to implement fault tolerance strategies, and the like.

Cellular module

[3328] In yet another example, a cellular module 37022 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize the performance and reliability of cellular network communication. Additional and/or alternative intelligence tasks may involve predictive tasks like estimating cellular network load during peak times, detection and identification tasks to monitor signal quality or interference, automation tasks to adjust handover mechanisms between cells, decision support tasks to improve frequency allocation, configuration tasks to enhance network scalability, and the like.

WiFi module

[3329] In yet another example, a WiFi module 37024 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize wireless local area network (WLAN) connectivity and performance. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting WiFi congestion in high-density areas, detection and identification tasks to analyze signal strength or interference sources, automation tasks to dynamically adjust channel allocations or access point settings, decision support tasks to enhance network configurations for specific environments, configuration tasks to enforce security protocols, and the like.

ORAN module

[3330] In yet another example, an ORAN (Open Radio Access Network) module 37026 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize and manage ORANs. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting traffic loads across base stations, detection and identification tasks to monitor ORAN performance, automation tasks to adjust radio resource allocations dynamically, decision support tasks to enhance multi-vendor ORAN configurations, configuration tasks to optimize spectrum usage, and the like.

Bluetooth module

[3331] In yet another example, a Bluetooth module 37028 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize short-range wireless communication for connected

devices. Additional and/or alternative intelligence tasks may involve predictive tasks like anticipating device pairing or traffic patterns, detection and identification tasks to analyze signal interference or connectivity issues, automation tasks to adjust device priorities dynamically, decision support tasks to enhance Bluetooth mesh configurations, configuration tasks to manage device security, and the like.

Satellite module

[3332] In yet another example, a satellite module 37030 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize satellite communication. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting weather-related signal degradation, detection and identification tasks to monitor satellite link performance, automation tasks to reroute traffic dynamically between satellites, decision support tasks to recommend satellite placement strategies, configuration tasks to optimize bandwidth utilization, and the like.

IoT network module

[3333] In yet another example, an IoT network module 37032 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to manage and optimize connectivity for Internet of Things (IoT) devices. Additional and/or alternative intelligence tasks may involve predictive tasks like estimating device activity patterns, detection and identification tasks to flag malfunctioning or compromised devices, automation tasks to dynamically allocate network resources for IoT traffic, decision support tasks to improve IoT gateway configurations, configuration tasks to enforce security measures, and the like.

Mesh network module

[3334] In yet another example, a mesh network module 37034 of the network layer 37000 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize and manage a set of decentralized and/or mesh networks. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting node connectivity or load imbalances, detection and identification tasks to locate failing or underperforming nodes, automation tasks to reroute traffic dynamically within the mesh, decision support tasks to improve network scalability and resilience, configuration tasks to optimize node configurations, and the like.

Social network module

[3335] In yet another example, a social network module 37036 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize user engagement, content delivery, and network connectivity across a social network. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting user activity patterns or viral content trends, detection and identification tasks to flag illegal content, automation tasks to curate personalized content feeds dynamically, decision support tasks to recommend moderation or community management strategies, configuration tasks to enforce privacy and security policies, and the like.

Power grid module

[3336] In yet another example, a power grid module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize the distribution, monitoring, and management of electricity across a set of power grids. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting power demand and identifying potential outages, detection and identification tasks to monitor grid stability or detect faults, automation tasks to dynamically reroute power to prevent overloads, decision support tasks to recommend maintenance schedules or grid expansion strategies, configuration tasks to implement energy storage or renewable integration policies, and many others.

Delivery network module

[3337] In yet another example, a delivery network module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize last-mile logistics, package tracking, and delivery efficiency across a delivery network. Additional and/or alternative intelligence tasks may involve predictive tasks like estimating delivery times and forecasting traffic patterns, detection and identification tasks to monitor delivery route bottlenecks or failed deliveries, automation tasks to dynamically reroute delivery vehicles for improved efficiency, decision support tasks to enhance fleet management strategies, configuration tasks to prioritize high-value or time-sensitive deliveries, and the like.

IoT-in-a-Box Module

[3338] In yet another example, an IoT-in-a-box module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to simplify the deployment, management, and integration of IoT-in-a-box systems within a network. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting device connectivity needs or resource consumption, detection and identification tasks to monitor device health or security vulnerabilities, automation tasks to configure and optimize IoT communication protocols, decision support tasks to enhance IoT deployment strategies, configuration tasks to standardize IoT device policies, and the like.

Industrial IT/OT module

[3339] In yet another example, an industrial IT/OT module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to bridge and optimize operations technology (OT) and/or information technology (IT) systems within industrial environments. Additional and/or alternative intelligence tasks may involve predictive tasks like forecasting equipment maintenance needs, detection and identification tasks to monitor system performance or potential failures, automation tasks to synchronize IT and/or OT workflows, decision support tasks to enhance production planning strategies, configuration tasks to ensure secure integration between IT and OT systems, and the like.

Adaptive networking and QoS platform

[3340] At the network layer 37000, AI capabilities for adaptive networking enable a platform for providing integrated, adaptive networking and quality of service (QoS) ("the adaptive networking

and QoS platform") for various enterprise and/or operational environments (in some implementations having various sets of facilities and/or environments contained therein). For example, enterprise and/or operational environments may include industrial environments, manufacturing environments, warehouse environments, shipping environments (e.g., container ports, shipping containers, shipyards, and the like), energy environments (e.g., fossil fuel energy environments such as oil rigs, pipelines, and refineries or nuclear energy environments such as nuclear power plants), retail environments, transportation environments (e.g., airports, train stations, bus stops, airplanes, train cars, and the like), transaction and banking environments, datacenter environments, compute environments, hospitality environments, gaming environments, and entertainment environments, among many others.

[3341] The adaptive networking and QoS platform may include, without limitation, caching of AI content and generative AI content (e.g., pre-generated instructions for operating machinery, customer risk assessments, pre-generated customer support responses, pre-generated explanations of medical conditions, optimal delivery routes for recurring logistics paths, path optimization results for robot pickers in commonly used zones, workload distribution recommendations, predictions for server failure risks, customer preference data, product recommendation results for repeat customers, pre-generated personalized marketing emails, and pre-generated maintenance instructions or escalation workflows for hardware issues, among many others); caching of AI models (e.g., predictive AI models such as predictive maintenance models for specific machinery, recommendation models, optimization models such as for optimizing factory floor layouts, credit risk assessment models, fraud detection models, load forecasting models, anomaly detection models, diagnostic models, inventory tracking models, pathfinding models, workload balancing models, pricing optimization models, resource allocation models, and the like) and generative AI models (such as for efficient, localized generation of AI content based on localized, contextual inputs and/or for localized, customized use); handling via edge capabilities of transaction events (including point-of-sale transactions, fraud detection, payment processing, energy trading, toll collection, bet placement, payouts, buy-ins, room booking, food and beverage purchases, event ticket purchases, loyalty and rewards transactions (e.g., points earning, points redemption, and membership upgrades), promotional transactions, reporting, payment resolution, and other transaction functions); execution by edge capabilities of gaming engine functions (such as enabling efficient, localized generation of AI-driven environments (e.g., large-scale foundation models for training intelligent agents), simulations, characters or personnel (including physical robots as well as software-based personnel for customer-facing interactions and/or other operations, as well as AI-based attributes, behaviors, and interactions for virtual or augmented systems, including AR/VR-enabled simulations, training modules, operational overlays, and the like); execution by edge capabilities of experience orchestration (such as generating immersive and interactive experiences tailored to localized context, operational needs, and/or personalization of generative content based on real-time understanding of target cohorts, operational conditions, or individual user preferences across various environments, including AR/VR simulations, dynamic visualizations, and situationally adaptive content delivery); and/or the like.

[3342] In embodiments, edge capabilities may be provided in edge systems, devices, switches, chipsets and other embodiments. Edge capabilities may include converged capabilities where a set of artificial intelligence capabilities, such as generative or other AI capabilities, are integrated with other networking capabilities, such as by providing them as part of the same device or system, providing them with shared or coordinated input/output channels, providing them with coordinated access to the same data, or the like.

Adaptive network capabilities

[3343] In embodiments, the adaptive networking and QoS platform includes caching of AI models for operations and use cases for various environments. In embodiments, the foregoing capabilities are enabled by adaptive networking capabilities as described in this disclosure and the documents incorporated by reference herein, such as by adapting a networking capability based on one or more of network conditions, based on a set of prioritization parameters, based on the content to be generated by a sender (e.g., an edge device) and or delivered over the network from a sender, based on the capabilities of a receiver, and/or based on the nature of usage by a receiver, among others. Such adaptive networking capabilities, based on such factors, may include, without limitation: a) adaptive selection of a networking protocol (e.g., a physical layer, transport layer, media access control layer, or other networking protocol); b) adaptive generation of content by an edge or other networking device that is capable of content generation, such as by generation of content of a suitable format, file type, size (e.g., by adaptive chunking of generative content into suitable block sizes or the like); c) adaptive storage of data on a network storage device, including caching for rapid access and/or long term storage; d) adaptive routing of network traffic, including based on QoS or channel conditions, cost of routing, and other factors; e) adaptive network coding, including setting of random linear network coding or other coding parameters; f) adaptive block sizing; g) adaptive data rate (including dynamically adjusting data rate according to a linear, convex, or concave adjustment curve); h) adaptive error correction, including adaptive setting of forward error correction parameters; i) adaptive session management, including setting timing of session initiative, termination and duration; j) adaptive and/or dynamic spectrum access management, including slicing usage available cellular or other spectrum based on any of the factors noted above; k) adaptive setting of mesh networking parameters; l) adaptive setting of software filters for repeaters, switches, routers and other software-defined transmission devices; m) adaptive configuration of mobile ad hoc network parameters; and others. Any one or more of these capabilities may be provided in an adaptive networking device or system that further includes a set of other AI capabilities, including generative AI capabilities for generation of content that is to be handled by the converged device or system.

[3344] In embodiments, the network layer 37000 may be configured to facilitate communications among smart network-capable systems, their users, and the like while leveraging AI to dynamically adapt network configurations and optimize data flow. The integration of edge computing with cloud services ensures that data processing and transaction and workflow orchestration can occur closer to the data source, reducing latency and enhancing real-time decision-making capabilities. AI algorithms may be capable of intelligently routing data (e.g., IoT data, sensor data, edge data,

transaction data, behavior data, social data, web data, crowdsourced data, media data, vector data, distributed data, and the like) through the most efficient network paths, whether they be cellular, WiFi, Open Radio Access Network (ORAN), Bluetooth, IoT communication protocols, and other communication protocols. For implementation, the network modules 37002 may utilize AI for predictive network traffic management, ensuring bandwidth is allocated where it is needed most, and cryptographic techniques for securing data at the edge. AI-driven anomaly detection systems may be employed to monitor network health and preemptively address potential disruptions.

[3345] In embodiments, the network layer 37000 may provide advanced computing capabilities at the edge of the network to process and manage transactions locally to improve response times, conserve bandwidth, enhance security, and the like. The network layer 37000 may be supported by edge devices, edge data centers, edge computing platforms, and the like. In embodiments, blockchain may be used to handle distributed transactions across multiple edge devices, offering security and transparency of transactions. AI may be used at the edge to make intelligent decisions based on the data processed, which may be particularly useful for personalized user experiences.

[3346] At the network layer 37000, caching of generative AI may be used for reducing the computational resources and energy consumption required for AI workloads, benefits that may optimize down to the resource layer 39000. In implementations, caching of AI workloads may involve output caching, feature caching, model state caching, and the like. Output caching may refer to storing outputs previously generated by the AI model. For example, output caching may involve caching previously generated output. Feature caching may refer to input data that has been pre-processed and transformed into a format suitable for model training or inference. Model state caching may refer to caching the state of the model at different points during its operation. Outputs, features, model states, and the like may be stored in a readily accessible storage layer, allowing for quick retrieval of information. In implementations, caching of AI workloads at the network layer 37000 may be enabled by a cache management system.

[3347] In embodiments, network layer 37000 may provide a content delivery network (CDN) for enabling the delivery of generative AI content. The integration of generative AI with CDNs can significantly enhance the functionality of CDNs with respect to content optimization, personalization, real-time content generation, and the like. For example, generative AI models may be configured to utilize user preferences, user behavior, and others to generate personalized content in real-time at the edge and/or demographic data, location data, and others to generate targeted content in real-time at the edge. In embodiments, each CDN may be configured as a network of distributed servers that are designed to deliver content. The CDNs may support caching of generative AI content, serving cached content from edge servers closest to users. The CDNs may be configured to perform load balancing during periods of high traffic. In embodiments, the CDNs may restrict access to content based on geographic locations in compliance with state laws and regulations. In some embodiments, the CDNs may use token authentication to control access to content (e.g., validating user's request to access content based on time restrictions, IP address limitations, the number of times content can be accessed, and the like).

[3348] In embodiments, content distribution may be significantly enhanced through the use of smart network-capable devices. Smart network-capable devices may include smart sensors, smart products, connected robotics, wearables, programmable logic controllers (PLCs), 3D printers, smart conveyors, smart containers, smart shelving systems, drones, distributed energy resources (DERs), smart machines, smart cooling systems, smart point-of-sale (POS) systems, smart TVs, and smart streaming devices, among many others. In embodiments, the smart network-capable devices can leverage AI functionalities for operations by analyzing user preferences, user behavior, user interactions, historical data, and the like. In embodiments, generative AI models may modify or generate content based on viewing conditions or user feedback. For example, AI can dynamically adjust video quality based on available bandwidth. In embodiments, the smart network-capable devices may use geolocation data to enforce geographic restrictions on activities, and the like.

AI model-driven gaming engine edge system

[3349] In embodiments, an AI model integrated with or supporting a gaming engine or other gaming system may be integrated with edge networking capabilities, such as for caching relevant content and/or producing generative AI content relevant to efficient, localized, personalized and/or or customized digital twin, virtual reality, augmented reality, metaverse or similar content experiences. This may include various edge networking device embodiments, such as by integrating the generative AI and gaming capabilities into a network box, router, switch, edge device, IoT system, mesh networking node, networking chip or chipset, networking system on chip (SoC), FPGA, or the like. Collectively these may be referred to herein, for simplified reference, as an AI model-driven gaming engine edge system.

[3350] In embodiments, the AI model-driven gaming engine edge system may be integrated with financial infrastructure systems, such as to enable embedded marketplaces and transactions, as described elsewhere in this disclosure and the documents incorporated by reference herein.

[3351] In embodiments, the AI model-driven gaming engine edge system may be configured to integrate with data handling systems that are capable of sensor and data fusion, i.e., receiving and processing content that provides an understanding of the context in which the information handled by the AI model-driven gaming engine edge system is used. This may include sensor data that provides information about a set of users (e.g., data detected from wearable devices or physiological data, data from IoT devices or cameras indicating information about the user or the user's environment, or the like); behavioral data about the user (e.g., clickstream data indicating the user's navigation patterns, transactional data, user purchases of content or other items, and other sources); demographic, psychographic and geographic data about populations users (including as determined by various similarity and clustering algorithms) and the like; and/or user-generated content data, including explicit indications of preferences, such as ratings, collaborative filtering and surveys and indirect indications, such as tastes, preferences and styles indicated by user-generated content, such as social media content (e.g., showing a style of dress or preference for a genre of music or content), written content, such as reflecting interests in particular content and/or indicators of user personality, and user interaction content, such as user interactions within

interactive game play (e.g., user dialog, navigation and other choices), with digital twins, and/or within virtual reality, augmented reality and/or metaverse environments. In embodiments, data and sensor fusion provide a holistic understanding of user content that can be used to train AI models (either at the edge or in the cloud) to have an accurate understanding of a user's position, duties, status, condition, preferences, personality, or the like. Once trained, the AI models can be deployed in the AI model-driven gaming engine edge system to generate content that is highly tuned to the context and preferences of the user.

[3352] In embodiments, the AI model-driven gaming engine edge system may be configured and trained to understand a user profile across different platforms, such as traditional linear media, enterprise platforms, entertainment platforms (e.g., music, film, television, streaming, and the like), and/or interactive platforms (e.g., sports betting, gaming, digital twin, social media, metaverse, virtual reality, augmented reality, and many others). Once the profile is understood, a resulting set of models may be deployed in the AI model-driven gaming engine edge system to provide highly tuned content and experiences to the user across all such channels. Such AI models may include long-term models that provide a general understanding of stable characteristics of the user (e.g., personality, long-term interests, relationships with other users, affinity for subjects, relationships with organizations) and short-term models that provide an understanding of immediate context and interests (e.g., content that is highly relevant to the user's current location).

[3353] In embodiments, different forms of content may be cached for efficient delivery to the user at the AI model-driven gaming engine edge system, and they may be presented in a set of integrated content-access interfaces that make different content forms available. As an understanding of an individual user (stable/long-term and/or immediate/contextual) is embedded in a set of AI models as noted above (aided by data and sensor fusion), that understanding can be linked to suites of content of any and all types, which can be gathered and cached by the AI model-driven gaming engine edge system and used to organize suites of content for presentation to the user, either in a unified "home screen" type interface or in forms where links to any one type of content are embedded into the other forms of content. Once content is gathered, it may be used both to configure content offerings and to generate, locally at the AI model-driven gaming engine edge system, generative content, such as text, images, videos, audio, offers, promotions, and many others. The AI model-driven gaming edge system may thus serve as an important enabling component for a broader platform as described herein that links various multiplatform content experiences, models that are trained on a holistic understanding of the user, including based on cross-platform interactions, are deployed at the edge to configure offerings and to generate relevant content for the user.

[3354] As implied by the above, the AI model-driven gaming engine edge system may include a set of distinct models, a set of hybrid models, or the like. For example, one AI model (or component) may be used to understand a long-term profile of a user, another AI model may be used to understand a short-term context of a user, another AI model may be used to discover, gather and/or configure a set of content for the user, other models may be used to generate highly personalized or customized generative content (such as involving text, voice, music, images,

videos, offers, promotions, or mixed modes of the above), and other models may be used to resolve transactions. In one such configuration, a set of models is used to predict a set of immediate preferences of the user and another set of models is used to generate a configuration of content or other content that is designed to satisfy those preferences.

[3355] In embodiments, an instance of an AI model-driven gaming engine edge system may be configured, provisioned and deployed for a single specific user, a defined set of similar users, or a defined specific group of users (such as a family or workgroup). The system may then be trained over time to learn the profile of the user or set, to gather and configure content and other content, and/or to generate highly relevant generative works.

[3356] In embodiments, the AI model-driven gaming engine edge may be trained to learn the personality of a user by observation of the user's cross-platform behavior (including various aspects noted in the discussion of sensor and data fusion above). By observing behavior (e.g., by machine vision), transactions, content consumption, interactions, and generated content, as well as by driving interactions that are designed to elicit indicators of personality (such as collaborative filtering, surveys and dialog), the AI model can classify the user according to psychometric or neurometric measures of personality (e.g., a Jungian, Myers-Briggs, or Keirsey Temperament type, other the like). These can be augmented by other indicators, such as social media and writing content by the individual, content preferences, and the like. With an understanding of personality embedded into an AI model, the personality can be used as an input to a content discovery and/or configuration system and/or a generative AI system to produce content that is appealing to the personality.

[3357] By integrating these technical features, the network layer 37000 is capable of supporting the complex networking requirements of various operational environments.

Data Layer

[3358] In embodiments, an artificial intelligence-enabled system of systems can have a data layer system of systems, a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to operate on fused data that can be the output of a sensor and/or data fusion system.

[3359] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include sensor data.

[3360] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include wearable device data.

[3361] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include social media data.

[3362] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include crowdsourced data.

[3363] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include website data.

[3364] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include distributed data.

[3365] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include API data.

[3366] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include edge data.

[3367] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include mobile data collector data.

[3368] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include industrial data.

[3369] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include transactional data.

[3370] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include transportation data.

[3371] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include value chain network data.

[3372] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include supply chain data.

[3373] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include demand data.

[3374] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include shipping data.

[3375] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include energy data.

[3376] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include enterprise data.

[3377] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include public data.

[3378] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include market data.

[3379] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include news data.

[3380] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include weather data.

[3381] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include contextual data.

[3382] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include health data.

[3383] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include demographic data.

[3384] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include insurance data.

[3385] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include simulation data.

[3386] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include experimental data.

[3387] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include synthetic data.

[3388] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include echoed data.

[3389] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include communication data.

[3390] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include behavioral data.

[3391] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include location data.

[3392] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include pricing data.

[3393] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include sales data.

[3394] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include tax data.

[3395] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include historical data.

[3396] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include physiological data.

[3397] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include environmental data.

[3398] In embodiments, the data layer system of systems having a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks can be configured to operate on fused data that can include key performance indicator (KPI) data.

[3399] In embodiments, an artificial intelligence-enabled system of systems 31000 can have a data layer system of systems 38000. With continuing reference to FIG. 276, the data layer system of systems 38000 can incorporate data modules 38002. The data layer system of systems 38000 can cooperate with a set of machine learning systems 38010, a set of artificial intelligence systems 38020, and/or a set of neural networks 38030 configured to operate on fused data 38050. By way of these examples, the fused data 38050 can be formed from the output of a sensor 38060 and/or a data fusion system 38070 each of which can be associated with one or more data sources 38500. The one or more data sources 38500 can include sensor data 38100, wearable device data 38110, social media data 38120, crowdsourced data 38130, website data 38140, and distributed data 38150. The one or more data sources 38500 can also include API data 38160 that can be accessed

through various APIs at 38550 in FIG. 277. With continuing reference to 189, the one or more data sources 38500 can also include edge data 38170, mobile data collector data 38180, and industrial data 38190. In further examples, one or more data sources 38500 can include transactional data 38200 in a financial infrastructure environment. In additional examples, one or more data sources 38500 can include transportation data 38210, which can be used to support transportation systems, self-driving vehicles, software-defined vehicles, and software-enabled vehicles. In yet more examples, one or more data sources 38500 can include value chain network data 38220, supply chain data 38230, demand data 38240, and/or shipping data 38250, which can be used to further the understanding of entity demand and supply management. In additional examples, one or more data sources 38500 can include energy data 38260, which can be used to further facilitate energy delivery, compute power, and data management at various edge applications. In further examples, one or more data sources 38500 can include enterprise data 38270, public data 38280, market data 38290, news data 38300, and weather data 38310. In additional examples, one or more data sources 38500 can include contextual data 38320 in various settings. In yet further examples, one or more data sources 38500 can include health data 38330 in support of various digital health and healthcare platforms. In additional examples, one or more data sources 38500 can include demographic data 38340, insurance data 38350, simulation data 38360, experimental data 38370, synthetic data 38380, echoed data 38390, communication data 38400, behavioral data 38410, location data 38420, pricing data 38430, sales data 38440, tax data 38450, historical data 38460, physiological data 38470, environmental data 38480, and key performance indicator data 38490.

[3400] FIG. 277 depicts additional exemplary features, capabilities, and interfaces of exemplary embodiments of a data layer system of systems 38000. In embodiments, the API data 38160 can be accessed by multiple APIs at 38550. By way of these examples, the APIs at 38550 can include one each of an ingestion 38552, parse 38554, analyze 38556, and elements that can be interconnected for providing intelligence-based and other derivatives of data sources, such as the multitude of data sources 38500. These elements may be controlled, configured and adjusted by a data layer control 38560.

[3401] In embodiments, data sources may be retrieved and/or received through the APIs 38550. In embodiments, the APIs may be an open/standardized APIs that, among other things, may equip the data layer system of systems 38000 for integration with and into current and emerging ecosystems. By way of these examples, APIs can further enable data layers to integrate into a wide range of data workflows and access the API data 38160 and various data sources 38500, such as corporate internal workflows, metadata for visual content, source data, inter-jurisdiction data workflows, digital rights management and the like.

[3402] In embodiments, the data layer system of systems 38000 can include, reference, and/or provide market orchestration elements 38570 that may facilitate use of data layer capabilities for various aspects of market orchestration, including, without limitations, software orchestrated transactions, software orchestrated marketplaces, and the like. Market orchestration elements 38570 may facilitate deployment of a data layer, such as a web service embodiment, as an integrated function of a market orchestration platform, such as an automated market orchestration

system of systems as described herein. In embodiments, the data layer system of systems 38000 may provide data and network pipeline capabilities for market orchestration.

[3403] The data layer system of systems 38000 can include reference and/or provide cross-market interaction capabilities 38580 that may enable leveraging various types of intelligence data layer, computation capabilities, and storage and data sourcing capabilities, as well as intelligence capabilities for cross-market interactions. Cross-market interaction capabilities 38580 may include interfaces to one or more marketplaces, transaction environments, and the like, so that, among other things, a data layer supportive of a system of systems (SoSs) may be configured with one market in a cross-market integration deployment as a source of data and with another market in the cross-market integration deployment as a consumer of the data layer. In embodiments, a similar arrangement may be constructed between two or more markets so that data in either market can be used as one or more data sources and can be influenced by data from another market. By way of these examples, cross-market interactions 38580 may be accomplished through one or more market-to-market data layers that form data pipelines for intelligent exchange of data among markets, such as data about buyers in one market and about sellers in another.

[3404] The data layer system of systems 38000 can include functions and processes 38590, for an exemplary market-oriented deployment may include software-oriented transaction functions and processes, automatic transaction transactions and processes, and the like. Functions and processes 38590 for the data layer system of systems 38000 may include signaling availability of data (e.g., emergence of an occurrence of source data) that impacts data produced by the data layer of the platform. Other exemplary functions and processes 38590 may include embedding into smart contracts, tokens, or other occurrence (e.g., an initiation of a smart contract and the like). Yet other functions and processes may include payments between/among machines and the like.

[3405] In embodiments, the data layer system of systems 38000 may include and/or be associated with data layer technology enablers 38600, such as 5G networking, artificial intelligence, visualization technology (e.g., VR/AR/XR), distributed ledger, and the like.

[3406] In embodiments, the data layer system of systems 38000 can include and/or leverage cloud-based virtualized containerization capabilities and services 38602, such as without limitation a container deployment and operation controller, such as Kubernetes 38604 and the like. Cloud-based virtualized containers 38602 provide for data layers to be deployed close to source data, thereby potentially reducing network bandwidth consumption or the potential for network disturbances in a data workflow, especially when addressing voluminous local situations such as in industrial and commercial environments.

[3407] In embodiments, technologies may be provided by and/or enabled by the data layer system of systems 38000 include intelligence services 38606, such as artificial intelligence, machine learning and the like. These intelligence services 38606 may be provided by the data layer SoSs 38000, or accessed (e.g., as third-party services) via one or more. The data layer control 38560 may be provided access to the intelligence services 38606. In embodiments, the data layer control 38560 may provide its own set of intelligence services.

[3408] In embodiments, the data layer system of systems 38000 may include transaction/market-oriented capabilities, and services that may include market platforms 38610, transaction flows 38612, users 38614, and content providers 38616. For multi-party transaction environments, the data layer SoSs 38000 may be configured and operated to satisfy a range of consumer needs for market analysis, transaction efficiencies, cost containment, buy/sell decisions and the like.

[3409] In embodiments, the data layer system of systems 38000 may include an exemplary intelligent data layer architecture 38650, as depicted in FIG. 278. The exemplary intelligent data layer architecture 38650 can include a controlled pipeline of data processing stages that process data from one of multiple data sources 38660. The controlled pipeline includes an ingestion stage 38662, an analysis stage 38664, a derived intelligence stage 38668, and a consumer visualization portal 38670. By way of these examples, the ingestion stage 38662 can receive and/or harvest data from one or more of the plurality of sources 38660. Processing at the ingestion stage 38662 may include parsing content of data sources, such as to determine structure, content, relationships among data elements, intended meaning of the data elements, relationships between data, structure, and meaning, and the like. In embodiments, an ingestion facility that may operate at the ingestion stage 38662 may be configured to be aware of data source aspects, such as structure, etc. Ingestion stage 38662 may be configured or adjusted by an operator of the layer, a platform controller, the data layer control 38560, and the like. Configuration of the ingestion stage 38662 may be based on one or more data structures that represent aspects of one or more of the data sources 38660. One such aspect is a location of the data source, such as an Internet or other type of address (e.g., URL, port number, stream identifier, publication and/or broad channel, sensor output location, and the like) from which data may be accessed, queried, pulled, downloaded, streamed or otherwise accessed. Further aspects of a data source may include interface methods and/or protocols, such as through data transfer handshakes, Internet Protocols, and/or query languages. Data block size, access rate (e.g., maximum, frequency, or other timing-related parameter related to accessing the data source), and the like may also be adjusted or configured. Yet other aspects of data source 38660 and the data layer architecture 38650 may include one or more meanings of data from the data source, such as may be represented through a data source ontology and the like. Information such as units, scalars and the like for numerical values may be provided. In an example of data sources providing measurement data, a first source may provide numerical values in inches, a second source may provide values in meters, and a third source may provide values in light years. This local data source context may prove useful for relating data sources. In an example of data sources providing reputation rating values, an ontology for each source that establishes a minimum value, maximum value, and adjustable increments between the two provide a way of establishing meaning of a data element from such a source.

[3410] In further examples, the data sources may impose or be arranged with a geometry/structure that may imbue meaning on data values, relationships among data values, and the like. Exemplary embodiments of structures that can impact the meaning and relationships of data value from a data source can be a hierarchical arrangement of the data.

[3411] In embodiments, the data layer system of systems 38000 can be configured to receive/retrieve and process data that is configured as a hierarchy and can be configured to assign a relationship attribute to a pair of data values that are configured as parent/child in the hierarchy. Likewise, a rule that may be applied in the hierarchy, such as certain types of changes to a parent data value impacting a corresponding child data value establish an immutable relationship between the data values as they are processed through the ingestion processing pipeline (e.g., ingestion, analysis, and intelligence).

[3412] In embodiments, automation can be introduced through programmatic configuration of data values in an ingestion stage execution data structure. These data values may be retrieved from, for example, an ingestion parameter portion of a data layer control data store 38672. Machine learning may facilitate training an artificial intelligence system to identify aspects of a data source that are relevant for configuring the ingestion stage 38662 to receive and process its data.

[3413] In embodiments, the data layer control 38560 may develop an understanding of data sources, such as meaning of data values. In embodiments, developing an understanding may be in context of an expected use of data from one or more data sources, such as use of the data for determining a status of a term of a smart contract, a result of a software orchestrated transaction, and the like. Further data from data sources may be understood within the context of other data sources. In an example of such understanding, a plurality of marketplace monitors may capture data regarding activity within a marketplace.

[3414] In embodiments, the data layer control 38560 may further be configured to maintain a schedule of collection activity for one or more of the data sources 38560. A collection schedule may be one of a plurality of aspects associated with ingestion that may be influenced by data sources and by data layer pipeline processing needs (e.g., to satisfy needs of a user of the data layer). Such a collection schedule may be based on a rate or occurrence of availability of new or revised data from a source. In embodiments, some data sources may produce new/updated data on a schedule determined from activities associated with the data source, such as a sample schedule of the sensor 38060, the data fusion 38062, and the like. In an example, a rule for a system that produces data available through a data source may limit data releases periodically (e.g., such as at the end of a work shift, once per day, and the like). In another example of data source-dependent collection activity may be motivated by events, such as completion of marketplace transactions and the like. In embodiments, the data layer control 38560 may monitor a port on a data network for an indication of data availability at a data source. When an indication is detected (e.g., a change in a data value of the port), an ingestion process may commence.

[3415] By way of these examples, other information of processes related to ingestion may include costs, such as costs to perform data source access, ingestion processing and the like. The cost for data collection may include access fees charged by a data source (e.g., subscription costs, access event costs, demand-based costs, and the like). Costs for data collection may be based at least in part on an amount that a consumer (e.g., a user of capabilities and output from a data layer) pays for access to information produced by the data layer that is based at least in part on data from the data source. By way of these examples, there may be a range of cost structures for source data

access, at least some of which may be based on data source reputation, relevance of data from the data source, timeliness of updates of the data from the data source, and the like. In embodiments, a data layer may access data from a data source and utilize it a plurality of times to produce layer intelligence for a plurality of users of the data layer. Costs for access and for the occurrences of use of the accessed data may be different from each other, such as a cost to access may be a multiple of (e.g., 2-time, 10-times and the like) of a cost for each subsequent occurrence of use of the accessed data.

[3416] In embodiments, the data layer system of systems 38000 may be configured as a producer of source data, so that a corresponding ingestion facility may be owned (and optionally operated) by the data producer. In an example of data source closely held within the data layer system of systems 38000, a data source may retain privacy of its source data by exposing, such as through publication and the like, an output of the owned data, which may include information derived from the source data or select portions of the source data, such as non-confidential information associated with marketplace transactions and the like.

[3417] In embodiments, further operation of the data layer architecture 38650 can include ingestion operation that may also be based on a method of data collection. In embodiments, the data sources 38660 may be part of a data supply chain. Exemplary embodiments of a data supply chain may include a physical chain, such as may be embodied by a set of physical sensors (e.g., industrial internet of things sensors) that capture physical activity (e.g., of an industrial machine, and the like) and provide a representation of that activity as a form of data. A physical connection, such as a set of networked devices (e.g., the Internet), may convey the representation of the activity produced by the sensor(s) to, for example, a physical access port (e.g., a networked computer and the like) from which a data layer may ingest this data. Other types of data collection may include logical supply chains, such as data marts, data marketplaces, aggregated data publishers, and the like. In embodiments, data representative of a physical activity, such as a production machine in an enterprise, may be provided through a physical interface that presents the data from a corresponding sensor as it changes in near real time. That same data may be provided through a logical interface, such as a data base that facilitates access to a plurality of values of data from the sensor, optionally with a capture time, capture sequence and the like to enable batched or delayed use of data from the data source.

[3418] In embodiments, the data layer control 38560 may execute in virtual containers on, for example, cloud computing systems that are distinct from a physical embodiment of the data layer control 38560.

[3419] In embodiments, the data layer architecture 38650 may include the analyze stage 38664 that may receive data from the ingestion stage 38662. The analyze stage 38664 may receive raw ingestion data, adapted ingestion data (e.g., contextually adjusted), data derived from ingestion data (e.g., differences between sequential accesses of a single data source), metadata associated with the ingestion data (e.g., validity window, units, access costs, and the like), and the like.

[3420] In embodiments, the analyze stage 38664 may perform various operations on ingestion stage parsing and other ingestion activity results based on a range of factors, such as comparing

data from a plurality of sources for similarity, fitness to a purpose, differences, based on types of data within or across data sources and the like. In embodiments, analysis may include comparing sources against a target use of intelligence derived from a data source. Analysis of ingestion results may attempt to determine if one or more data elements from a data source may meet consumption target requirements, such as meeting a validity time constraint, an accuracy constraint, a frequency of update constraint, relevance to a consumption subject matter focus, and the like. In embodiments, a data layer may target providing intelligence for buyers of services in a software orchestrated transaction marketplace. The data layer architecture 38650 may include functionality to publish or otherwise convey requests for data, such as types of data, and the like that one or more data sources may attempt to meet. The data layer control 38560 may determine if ingested data meets requirements of the published request for data, such as if the data complies with one or more parameters in the request.

[3421] In embodiments, the data layer control 38560 may facilitate configuring data in the layer for publication, such as configuring one or more advertisements that characterize the ingested data in terms of potential intelligence value, relevance and the like. Examples include making data, such as derived intelligence data available on a marketplace (e.g., configuring indexing schemes and the like), making the content searchable (e.g., identifying keywords, terms, values, or the like that may facilitate discovery of intelligence derived from the ingested data through use of a search capability. The data layer control 38560 may facilitate access visibility to information of the data layer by publishing, communicating, or broadcasting samples of the data over a network, directly to potential consumers and the like.

[3422] In embodiments, the analyze stage 38664 may suggest, predict, and/or estimate value of data for a plurality of different consumers. These estimates may be used by the data layer control to impact intelligent data layer functions, such as data layer intelligence pricing and the like that may be differentiated for different users. Further, such analysis may indicate that intelligence derived from a first data source may be more or less valuable to different target consumers.

[3423] In embodiments, the data layer control 38560 may use feedback from intelligent data layer users regarding, among other things, usefulness of intelligence derived from one or more data sources 38660 to facilitate ingestion and analysis activities and the like. In an example, positive feedback on intelligence derived from a data source may result in the data layer control 38560 to make use of the data source for deriving other types of intelligence and the like.

[3424] In embodiments, the data layer control 38560 may ingest data from a plurality of data sources; each such set of data may be individually analyzed. In embodiments, data from a plurality of data sources may be parsed, so that data with similar characteristics (e.g., data that is indicative of a buyer reputation) may be aggregated and analyzed. Examples of multiple data sources that may provide data with similar characteristics include mobile devices, types of sensors, media-focused transaction systems (e.g., content buying, advertising placement, rights management and the like). As noted above, the data layer control 38560 may communicate configuration data (e.g., sets of data that enable the analyze stage 38664 to perform various analysis functions). In

embodiments, the data layer control 38560 may provide a set of analysis algorithms that may execute on one or more processors.

[3425] In embodiments, the data layer architecture 38650 may include the intelligence stage 38668. By way of many examples, an intelligence stage 38668 may utilize artificial intelligence capabilities to develop an understanding about data sources including, among things, uses of data, values of data, applicability of data, collection patterns and relevance to intelligence consumption and the like. Additional intelligence that may be derived by intelligence stage may include, without limitation, layer specific data relevance, relevance of data from one layer to another, value of intelligence to a consumer, such as to a transactor. By way of many examples, intelligence stage may derive intelligence useful for forming new marketplaces from transactional data gathered from an existing marketplace.

[3426] In embodiments, the data layer architecture 38650 may include a consumer visualization portal 38670 that may communicate with elements of the data layer pipeline, such as the intelligence stage 38668 from which the consumer visualization portal may receive derived intelligence. In embodiments, the consumer visualization portal 38670 may facilitate access to derived intelligence (and optionally to other data of the data layer architecture 38650). The consumer visualization portal may announce availability of derived intelligence to a preconfigured set of consumers and candidate consumers through use of a messaging channel (e.g., SMS messaging and the like). The consumer visualization portal may announce derived intelligence through various other techniques including, broadcasting across one or more communication channels (e.g., TWITTER™, X™, and the like). The consumer visualization portal may deliver at least select derived intelligence to intelligent data layer consumers based on a subscription or similar arrangement between the consumer(s) and the data layer. In embodiments, the consumer visualization portal may reference publication configuration data that may identify which consumers are to receive which portion(s) of intelligence derived from which data source and cause the derived intelligence to be provided to (and/or made available to) one or more consumers based on this intelligence publication data. The consumer visualization portal may also communicate with the data layer control 38560, such as to receive configuration, access intelligence data, analyzed data, ingested data and the like.

[3427] In embodiments, the consumer visualization portal 38670 may further receive from one or more data layer data consumers, consumer preferences for interfacing with the consumer, requests for updates to previously communicated derived intelligence data, requests for onboarding, feedback on uses of derived intelligence data and the like. In embodiments, a consumer may communicate a derived intelligence delivery schedule to the consumer visualization portal where it may be combined with other intelligence delivery data, such as other consumer delivery schedules, and the like and utilized by the consumer visualization portal and/or the data layer control when performing derived intelligence delivery and communication functions.

[3428] In embodiments, the data layer control 38560 may provide configuration, control, storage, and processing capabilities, such as for providing access to algorithms from an algorithm portal 38680, facilitating access by the data layer control 38560 to intelligence services 38682, managing

storage of a data store, managing storage of intelligent data layer ingestion data and outcomes, analysis outcomes, derived intelligence and the like in a data store 38672, and without limitation providing a mechanism by which a user, such as an owner and/or operator of the data layer architecture 38650 can configure and otherwise interface with system of the data layer system of systems 38000.

[3429] In another exemplary embodiment, the data layer control 38560 may facilitate access to analysis algorithms by the analyze stage 38664. Further the data layer control 38560 may work cooperatively with an algorithm portal 38680 to receive algorithms for analysis, ingestion, deriving intelligence, and the like. By way of many examples, the data layer control 38560 and the data source 38500 may identify and/or provide one or more ingestion algorithms for performing ingestion actions on data provided. The algorithm may be provided through the algorithm portal 38680 received and optionally vetted by the data layer control 38560 and stored in the data store 38672. In another exemplary embodiment of use of the algorithm portal 38680, a consumer may provide an algorithm for deriving intelligence from data under the consumer's control, such as in a marketplace transaction environment in which a seller provides transaction data as a source of data to the data layer for processing, optionally with other relevant data, for deriving intelligence associated with seller marketplace activities. In embodiments, deployment of a data layer as part of a data workflow for an enterprise may involve adapting existing workflow steps with intelligent data layer capabilities. By way of many examples, a purchasing department of an enterprise may have a set of algorithms that are used to process sales forecast data for generating purchasing guidelines. A data layer may be constructed for the enterprise that produces intelligence regarding the generated purchasing guidelines by utilizing sales forecast processing algorithms that have been uploaded through, for example, the algorithm portal 38680.

[3430] In embodiments, the intelligence services 38682 may include a range of intelligence functions and capabilities including, without limitation artificial intelligence functions, machine learning functions, neural network functions, prediction capabilities, and many others. In an example of intelligence services 38682 for the data layer system of systems 38000, an ingestion stage 38662 may provide data from the data sources 38660, along with associated descriptive information (e.g., metadata, structural data, ontology data and the like) to a self-learning neural network capability of the intelligence services 38682 to aid in determining an approach to parsing the data source.

[3431] In embodiments, the intelligence services 38682 may further have access to subject matter associated intelligence, such as cross-market intelligence gathered through processing, optionally external to the data layer architecture 38650, marketplace configuration, operational, and transaction outcomes for different sets of cross-market offerings. By way of further examples using intelligence services for ingestion, this subject matter intelligence can be applied when a data source is determined to be related to a product or other offering that is similar to products or offerings on which the subject matter intelligence is based. In such instances when a source of data relates to a product (e.g., mobile device) and subject matter intelligence known to the intelligence services 38682 is based on or associated with mobile device technology, the corresponding

intelligence services may be utilized for enhancing/optimizing pipeline operations being performed on the source data.

[3432] In embodiments, the exemplary data layer architecture 38650 can include a user interface 38690 through which a user, such as an operator and the like, may interface with systems of the data layer and the like (e.g., query and maintain data in the data layer data store 38672). The user interface 38690 may facilitate configuring portions of the data layer, such as the algorithm portal, data retention rules for accessed by the data layer control 38560, prioritization of use of the data layer resources by data consumers, and the like.

[3433] The data layer system of systems 38000 may serve as a flexible and scalable infrastructure designed to enable the seamless integration, ingestion, and analysis of data from an extensive range of sources positioned within the artificial intelligence convergence system of systems 31000. The data layer system of systems 38000 can be positioned within the multi-layered architectures artificial intelligence convergence system of systems 31000 and can support the generation of actionable intelligence and foster integration across diverse ecosystems. The data layer system of systems 38000 supports a modular design to support corporate operations, market orchestration, and cross-functional workflows, as these systems of systems become important for organizations navigating both established and emerging technologies.

Sensor Data 38100

[3434] Sensor data may include real-time readings from devices monitoring environmental, industrial, or physical conditions. These sensors can provide valuable inputs for applications such as manufacturing quality control, environmental monitoring, and smart city management. For example, in an industrial setting, sensors might monitor machinery vibrations to predict maintenance needs. In environmental monitoring, air quality sensors may track pollution levels, enabling responsive public health measures. The data layer system of systems may process such high-frequency inputs, ensuring timely integration into decision-making frameworks.

Wearable Device Data 38110

[3435] Wearable devices may generate diverse datasets ranging from health metrics, such as heart rate and activity levels, to geolocation tracking. These devices often serve applications in healthcare, fitness, and workforce management. A healthcare provider, for instance, might leverage wearable data to monitor chronic conditions and deliver personalized care plans. The data layer system of systems can analyze this data in real-time, identifying trends and generating predictive insights to improve health outcomes.

Social Media Data 38120

[3436] Social media platforms may produce vast volumes of unstructured and structured data, including text, images, and videos. Businesses can use this data to gauge public sentiment, analyze trends, and optimize marketing strategies. For instance, a retail brand might monitor customer feedback on social platforms to refine its offerings. The data layer system of systems may enable sentiment analysis and predictive modeling to help businesses adapt to market demands effectively.

Crowdsourced Data 38130

[3437] Crowdsourced data may encompass information contributed by a distributed network of individuals, such as reviews, surveys, and collaborative maps. This data can be crucial for applications requiring diverse perspectives, such as urban planning or product development. For example, transportation authorities might use crowdsourced traffic reports to optimize routes and reduce congestion. The data layer system of systems can aggregate and contextualize such data to provide actionable insights.

Website Data 38140

[3438] Website data may include user interactions, clickstreams, and behavioral patterns. This data is invaluable for e-commerce, digital marketing, and user experience optimization. For instance, an online retailer might analyze browsing habits to recommend products, enhancing customer satisfaction and sales. The data layer system of systems can process large volumes of web traffic data, while its intelligence services may identify patterns to improve user engagement.

Distributed Data 38150

[3439] Distributed data may refer to datasets stored across multiple locations, such as cloud services or decentralized networks. This type of data is essential for global enterprises and blockchain applications. For example, a supply chain management system might rely on distributed data to track goods across international borders. The data layer system of systems may harmonize distributed datasets, ensuring consistency and enabling comprehensive analysis.

API Data 38160

[3440] API data may provide programmatic access to external systems, enabling the integration of third-party services and real-time data retrieval. Applications may include weather forecasting, financial market analysis, and social media monitoring. For example, a fintech platform might use API data to retrieve real-time stock prices and integrate them into portfolio management tools. The data layer system of systems may seamlessly ingest and process API data, supporting dynamic and responsive applications.

Edge Data 38170

[3441] Edge data may originate from devices operating at the periphery of a network, such as IoT devices in remote locations. This data is often used in scenarios requiring low-latency processing, such as autonomous vehicles or industrial automation. For instance, edge devices in a factory might monitor production line conditions and trigger immediate adjustments. The data layer system of systems may process edge data locally or in the cloud, enabling rapid responses and enhanced operational efficiency.

Mobile Data Collector Data 38180

[3442] Mobile data collectors may gather information from field operations, such as surveys, inspections, or environmental sampling. This data is vital for industries like agriculture, construction, and utilities. For example, agricultural researchers might use mobile devices to record soil conditions and crop health. The data layer system of systems may aggregate and analyze this data, providing actionable insights for field operations.

Industrial Data 38190

[3443] Industrial data may encompass metrics from manufacturing processes, such as temperature, pressure, and production rates. This data is critical for optimizing operations, ensuring quality, and predicting maintenance needs. For instance, a factory might use the data layer system of systems to monitor and analyze machine performance, reducing downtime and improving efficiency.

Transactional Data 38200

[3444] Transactional data may include records of financial or operational exchanges, such as purchases, sales, or payments. This data is foundational for industries like retail, finance, and logistics. A retail chain might use the data layer system of systems to analyze point-of-sale transactions, identifying trends and informing inventory strategies.

Transportation Data 38210

[3445] Transportation data may include vehicle telemetry, traffic patterns, and logistics tracking. This data is essential for optimizing supply chains and improving urban mobility. For example, a logistics provider might analyze transportation data to reduce delivery times and fuel consumption. The data layer system of systems may process this data in real-time, enhancing decision-making capabilities.

Value Chain Network Data 38220

[3446] Value chain network data may map relationships and dependencies across suppliers, manufacturers, and distributors. This data can support strategic planning and risk management. For instance, a manufacturer might use the data layer system of systems to visualize its supply chain and identify bottlenecks, improving overall efficiency.

Supply Chain Data 38230

[3447] Supply chain data may track the movement of goods, inventory levels, and supplier performance. This data is critical for ensuring timely deliveries and managing costs. A retailer might use the data layer system of systems to monitor supply chain performance and predict potential disruptions, enabling proactive measures.

Demand Data 38240

[3448] Demand data may capture customer preferences, market trends, and purchasing patterns. This data is vital for forecasting and aligning production with market needs. A consumer goods company might use the data layer system of systems to analyze demand data, optimizing product availability and marketing strategies.

Shipping Data 38250

[3449] Shipping data may include information about freight movements, carrier performance, and delivery timelines. This data is crucial for logistics and e-commerce. For instance, an online retailer might use the data layer system of systems to track shipments and provide real-time updates to customers.

Energy Data 38260

[3450] Energy data may capture metrics from utilities, such as consumption rates, production levels, and grid performance. This data supports sustainability initiatives and operational efficiency. A smart grid operator might use the data layer system of systems to balance supply and demand, reducing energy waste.

Enterprise Data 38270

[3451] Enterprise data may encompass internal records, such as HR files, financial statements, and operational metrics. This data is essential for decision-making and regulatory compliance. A corporation might use the data layer system of systems to integrate and analyze enterprise data, supporting strategic planning.

Public Data 38280

[3452] Public data may include datasets from government agencies, NGOs, and other public entities. This data can inform policy-making, research, and community initiatives. For example, a city planner might use the data layer system of systems to analyze public transportation data, improving service delivery.

Market Data 38290

[3453] Market data may capture financial metrics, stock prices, and commodity trends. This data is crucial for investment and trading strategies. A hedge fund might use the data layer system of systems to analyze market data, identifying opportunities and mitigating risks.

News Data 38300

[3454] News data may include headlines, articles, and broadcasts. This data can inform sentiment analysis, trend detection, and crisis management. A public relations firm might use the data layer system of systems to monitor news coverage, shaping communication strategies.

Weather Data 38310

[3455] Weather data may include forecasts, historical records, and real-time observations. This data supports agriculture, logistics, and disaster response. A logistics company might use the data layer system of systems to optimize routes based on weather conditions.

Contextual Data 38320

[3456] Contextual data may provide situational awareness by integrating multiple data sources. This data supports decision-making in dynamic environments. A retailer might use the data layer system of systems to combine sales data with demographic insights, tailoring marketing efforts.

Health Data 38330

[3457] Health data may include patient records, diagnostic results, and public health statistics. This data supports personalized medicine and population health management. A healthcare provider might use the data layer system of systems to analyze health data, improving care delivery.

Demographic Data 38340

[3458] Demographic data may capture information about populations, such as age, income, and education levels. This data informs marketing, policy-making, and resource allocation. A non-profit organization might use the data layer system of systems to target programs effectively.

Insurance Data 38350

[3459] Insurance data may include claims, risk assessments, and policy details. This data supports underwriting, fraud detection, and customer service. An insurance company might use the data layer system of systems to analyze claims data, reducing losses and improving customer satisfaction.

Simulation Data 38360

[3460] Simulation data may model hypothetical scenarios, such as financial projections or disaster response plans. For example, an organization planning for disasters from weather can continue its preparations by using simulation data processed through the data layer system of systems 38000. This data can allow the modeling of potential impacts and the development of optimized response strategies. A financial institution may also employ simulation data to project future market trends or assess the implications of new investment strategies. By integrating simulation outputs, the data layer system of systems can enable informed decision-making and facilitate the testing of complex hypotheses across diverse industries.

Experimental Data 38370

[3461] Experimental data may originate from controlled tests and research activities, capturing observations, measurements, and results. For instance, pharmaceutical companies may conduct clinical trials and generate experimental data related to drug efficacy. This data can then be processed and analyzed by the data layer system of systems 38000 to identify patterns and support regulatory submissions. Similarly, an automotive company might analyze experimental data from vehicle safety tests to enhance design standards. The flexibility of the data layer system allows for integration of experimental data into broader analytical frameworks, fostering innovation and development.

Synthetic Data 38380

[3462] Synthetic data may refer to artificially generated datasets that mimic real-world data patterns. These datasets are often used in scenarios where privacy concerns limit the availability of real data, such as training machine learning models. For example, a financial institution may generate synthetic transaction data to test fraud detection algorithms. The data layer system of systems 38000 can synthesize, store, and validate such datasets, enabling organizations to derive insights while adhering to privacy regulations. This capability is particularly valuable for advancing AI and machine learning applications.

Echoed Data 38390

[3463] Echoed data may involve data derived from repetitions or feedback loops, often used to validate or enhance existing datasets. For instance, a recommendation system may use echoed data to refine its predictions by analyzing user responses to previous suggestions. The data layer system of systems 38000 can seamlessly manage echoed data streams, incorporating them into real-time analytics to improve service delivery and user experience.

Communication Data 38400

[3464] Communication data may include emails, messages, and voice recordings, providing insights into interpersonal or organizational interactions. For example, businesses may analyze communication data to identify collaboration patterns and improve team productivity. The data layer system of systems 38000 can process such data while ensuring compliance with privacy and security standards. Applications in customer service, for instance, may benefit from sentiment analysis of customer interactions to enhance support experiences.

Behavioral Data 38410

[3465] Behavioral data may capture individual or group actions and preferences, often derived from online activities, app usage, or consumer interactions. A retail company might analyze behavioral data to personalize marketing campaigns or optimize product recommendations. The data layer system of systems 38000 can integrate behavioral data with other sources, enabling organizations to uncover trends and enhance customer engagement strategies.

Location Data 38420

[3466] Location data may include GPS coordinates, geofencing insights, and spatial movement patterns. This data can support applications in transportation, urban planning, and retail. For instance, logistics companies may use location data to optimize delivery routes, while urban planners may analyze it to improve public infrastructure. The data layer system of systems 38000 provides a robust framework for processing and integrating location data to support geospatial analytics and decision-making.

Pricing Data 38430

[3467] Pricing data may encompass historical price points, market comparisons, and promotional effects. Retailers may use this data to optimize pricing strategies, ensuring competitiveness and profitability. The data layer system of systems 38000 can process large volumes of pricing data, integrating it with demand and market data to enhance predictive pricing models.

Sales Data 38440

[3468] Sales data may include transaction records, customer purchase histories, and revenue trends. Businesses can analyze sales data to identify growth opportunities and optimize inventory management. For instance, a retailer might use the data layer system of systems 38000 to track sales performance across regions and refine marketing campaigns accordingly.

Tax Data 38450

[3469] Tax data may involve records of tax filings, liabilities, and regulatory compliance. Corporations can use this data to ensure accuracy in financial reporting and minimize audit risks. The data layer system of systems 38000 may facilitate the integration of tax data with financial and enterprise data, streamlining compliance processes.

Historical Data 38460

[3470] Historical data may provide context by capturing trends and events from the past. This data is invaluable for forecasting, policy-making, and research. A financial institution might use historical market data to identify cycles and predict future trends. The data layer system of systems 38000 can enable efficient storage and retrieval of historical data, enhancing long-term analysis and strategic planning.

Physiological Data 38470

[3471] Physiological data may include metrics such as heart rate, blood pressure, and respiratory rate, often collected through medical devices or wearables. Healthcare providers may use this data for personalized treatment plans and remote monitoring. The data layer system of systems 38000 can process physiological data in real-time, supporting advanced health analytics and early detection of medical conditions.

Environmental Data 38480

[3472] Environmental data may encompass metrics related to air quality, temperature, and water levels, often used in sustainability and disaster response efforts. For example, environmental monitoring organizations might analyze this data to predict and mitigate climate impacts. The data layer system of systems 38000 can integrate environmental data with simulation and public datasets to support comprehensive environmental management strategies.

Key Performance Indicator Data 38490

[3473] Key performance indicator KPI data may track metrics aligned with organizational objectives, such as customer satisfaction scores or operational efficiency rates. Businesses can use this data to evaluate performance and inform strategic initiatives. The data layer system of systems 38000 may enable real-time monitoring and analysis of KPIs, ensuring that organizations remain agile and focused on their goals.

Resource Layer

Resource modules

[3474] Referring to FIG. 271 and FIG. 279, the AI convergence system of systems 31000 may comprise a resource layer system of systems ("resource layer") 39000 having a set of resource modules 39002. In embodiments, the set of resource modules 39002 may include an advanced compute resources module 39004, a quantum computing resources module 39006, an edge computing resources module 39008, an intelligent agent resources module 39010, a model resources module 39012, an expert resources module 39014, an attention resources module 39016, a sensor resources module 39018, a robotic resources module 39020, a natural resources module 39022, a food resources module 39024, a geological resources module 39026, a physical storage resources module 39028, a physical security resources module 39030, a human resources module 39032, a machine resources module 39034, an advertising resources module 39036, a communication resources module 39038, a cloud-based resources module 39040, an enterprise resources module 39042, a facility resources module 39044, a heating resources module 39046, a cooling resources module 39048, an equipment resources module 39050, an information technology (IT) resources module 39052, a networking resources module 39054, and/or a material resources module 39056. In some embodiments, the resource modules 39002 of the resource layer 39000 may further include an energy resources module, a distributed energy resources module, and/or a renewable energy resources module. In some embodiments, the resource modules 39002 of the resource layer 39000 may further include a healthcare resources module. In implementations, the resource modules 39002 of the resource layer 39000 may include an industrial resources module. In embodiments, the resource modules 39002 may include a financial resources module and/or a transactional resources module. In some implementations, the resource modules 39002 may further include a transportation resources module, a physical delivery resources module, and/or a vehicular resources module. In some embodiments, the resource modules 39002 of the resource layer 39000 may further comprise a shipping container resources module, a shipping resources module, and/or an additive manufacturing resources module.

[3475] In embodiments, the resource layer 39000 enables the automated and optimized provisioning of resources used to support various systems and services of an enterprise and/or

operational environment. For example, operational environments may include industrial environments, manufacturing environments, warehouse environments, shipping environments (e.g., container ports, shipping containers, shipyards, and the like), energy environments (e.g., fossil fuel energy environments such as oil rigs, pipelines, and refineries or nuclear energy environments such as nuclear power plants), retail environments, transportation environments (e.g., airports, train stations, bus stops, airplanes, train cars, and the like), research environments, transaction and banking environments, datacenter environments, compute environments, hospitality environments, gaming environments, and entertainment environments, among many others.

[3476] Each resource module of the resource layer 39000 may include one or more processors and one or more memories storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize the provisioning of the set of resources corresponding to that particular module. The machine learning systems may be trained using supervised, unsupervised, or reinforcement learning techniques on historical resource allocation data, usage patterns, performance metrics, and simulated data (e.g., data generated from a simulation engine, a foundation world model, a set of digital twins, or the like). The neural networks may comprise multiple layers including, but not limited to, convolutional layers, recurrent layers, and fully connected layers, wherein the layers are configured to process input data and generate output predictions for resource optimization. The artificial intelligence systems may implement various algorithms including decision trees, random forests, and gradient boosting machines to analyze resource utilization patterns and make real-time optimization decisions. The resource layer 39000 may further include a simulation engine or set of foundation world models operating in parallel with the machine learning systems, artificial intelligence systems, and/or neural networks, wherein the simulation engine and/or set of foundation world models are configured to generate real-time simulations of resource allocation scenarios, validate optimization decisions, and provide feedback to the machine learning models for continuous improvement and adaptation to changing conditions. In embodiments, the simulation engine may employ physics-based modeling, discrete event simulation, or agent-based modeling techniques to accurately represent the dynamics of resource utilization and system behavior. In some embodiments, each resource module of the resource layer 39000 may include one or more processors and one or more memories storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to perform an intelligence task related to that particular resource module. Such intelligence tasks may include other types of optimizations, as well as predictive intelligence tasks, detection and/or identification intelligence tasks, decision support tasks, automation tasks, configuration and/or control intelligence tasks, and the like.

Advanced compute resource module

[3477] For example, the advanced compute resource module 39004 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of advanced compute resources. In embodiments, advanced compute resources may include high-performance computing (HPC) systems, graphics processing units (GPUs), tensor processing units (TPUs), field-programmable gate arrays (FPGAs), AI accelerators, distributed computing systems, and the like.

Quantum computing resources module

[3478] In another example, the quantum computing resources module 39006 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of quantum computing resources. In embodiments, quantum computing resources may include quantum processing units (QPUs), quantum annealers, gate-based quantum processors, quantum error correction systems, and the like.

Edge computing resources module

[3479] In another example, the edge computing resources module 39008 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of edge computing resources. In embodiments, edge computing resources may include Internet of Things (IoT) gateways, edge servers, edge AI accelerators, real-time analytics platforms, and the like.

Intelligent agent resources module

[3480] In another example, the intelligent agent resources module 39010 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of intelligent agent resources. In embodiments, intelligent agent resources may include autonomous agents, chatbot frameworks, multi-agent systems, decision-making algorithms, and the like.

Model resources module

[3481] In another example, the model resources module 39012 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of model resources. In embodiments, model resources may include predictive models, prescriptive analytics models, simulation models, and the like.

Expert resources module

[3482] In another example, the expert resources module 39014 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of expert resources. In embodiments, expert resources may include rule-based systems, decision-support systems, expert knowledge bases, and the like.

Attention resources module

[3483] In another example, the attention resources module 39016 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of attention resources. In embodiments, attention resources may refer to the human, computational, and/or AI-driven focus allocated to specific tasks, processes, or decision-making. For example, AI models may be trained to focus on anomalies in sensor data or prioritize alerts in high-traffic environments. In embodiments, attention resources may be hybrid attention resources, such as a combination of human and AI focus, where AI filters or highlights critical information for human review.

Sensor resources module

[3484] In another example, the sensor resources module 39018 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of sensor resources. In embodiments, sensor resources may include environmental sensors, motion sensors, imaging sensors, health sensors, and the like.

Robotic resources module

[3485] In another example, the robotic resources module 39020 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of robotic resources. In embodiments, robotic resources may include industrial robots, collaborative robots (cobots), autonomous mobile robots (AMRs), material handling robots, pick-and-place robots, robotic arms, robotic dock workers, maintenance robots, hospitality robots, cleaning robots, pipeline inspection robots, baggage handling robots, security robots, cargo loading robots, and the like.

Natural resources module

[3486] In another example, the natural resources module 39022 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of natural resources. In embodiments, natural resources may include water resources, mineral resources, forest resources, and the like.

Food resources module

[3487] In another example, the food resources module 39024 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of food resources. In embodiments, food resources may include agricultural produce, processed foods, perishable items, and the like.

Geological resources module

[3488] In another example, the geological resources module 39026 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of geological resources. In embodiments, geological resources may include minerals, fossil fuels, underground water reserves, and the like.

Physical storage resources module

[3489] In another example, the physical storage resources module 39028 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of physical storage resources. In embodiments, physical storage resources may include warehouses, cold storage facilities, data storage units, and the like.

Physical security resources module

[3490] In another example, the physical security resources module 39030 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of physical security resources. In embodiments, physical security resources may include surveillance systems, access control systems, alarm systems, locking systems, and the like.

Human resources module

[3491] In another example, the human resources module 39032 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of human resources.

Machine resources module

[3492] In another example, the machine resources module 39034 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of machine resources.

Advertising resources module

[3493] In another example, the advertising resources module 39036 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of advertising resources. In embodiments, advertising resources may include digital ad placements, social media campaigns, targeted marketing systems, and the like.

Communication resources module

[3494] In another example, the communication resources module 39038 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of communication resources. In embodiments, communication resources may include email systems, messaging platforms, telecommunication networks, and the like.

Cloud-based resources module

[3495] In another example, the cloud-based resources module 39040 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of cloud-based resources. In embodiments, cloud-based resources may include virtual machines, containerized services, distributed databases, and the like.

Enterprise resources module

[3496] In another example, the enterprise resources module 39042 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks

configured to optimize, by computing hardware, the provisioning of the set of enterprise resources. In embodiments, enterprise resources may include financial systems, operational workflows, customer relationship management (CRM) platforms, and the like.

Facility resources module

[3497] In another example, the facility resources module 39044 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by facility management systems, the provisioning of the set of facility resources. In embodiments, facility resources may include factories, textile mills, assembly plants, manufacturing facilities, distribution centers, fulfillment centers, storage facilities, warehouse facilities, container ports and terminals, shipyards, freight forwarding facilities, oil rigs and offshore platforms, refineries, natural gas processing plants, pipeline pump stations, coal mining facilities, nuclear power plants, uranium enrichment facilities, retail facilities, airport facilities, bus terminal facilities, platforms and waiting areas, ticketing facilities, maintenance facilities, banking facilities, ATMs and kiosks, call center facilities, office facilities, datacenters, high-performance computing (HPC) facilities, cloud computing facilities, hospitality facilities, and entertainment facilities, among many others.

Heating resources module

[3498] In another example, the heating resources module 39046 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of heating resources. In embodiments, heating resources may include boilers, geothermal heating systems, district heating networks, and the like.

Cooling resources module

[3499] In another example, the cooling resources module 39048 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of cooling resources. In embodiments, cooling resources may include HVAC systems, liquid cooling systems, evaporative cooling units, and the like.

Equipment resources module

[3500] In another example, the equipment resources module 39050 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of equipment resources. In embodiments, equipment resources may include industrial tools, construction equipment, diagnostic machines, and the like.

Information technology resources module

[3501] In another example, the information technology (IT) resources module 39052 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of IT resources. In embodiments, IT resources may include servers, data management systems, and the like.

Networking resources module

[3502] In another example, the networking resources module 39054 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of networking resources. In embodiments, networking resources may include switches, routers, hubs, bridges, network interface cards (NICs), access points (APs), modems, firewalls, network management software, virtual private networks (VPNs), network operating systems, software-defined networking (SDN) platforms, load balancers, wireless networking resources (e.g., Wi-Fi networks, cellular networks, and satellite networks), edge networking devices, IoT gateways, content delivery networks (CDNs), traffic analyzers, and many others.

Material resources module

[3503] In another example, the material resources module 39056 may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of material resources. In embodiments, material resources may include raw materials (e.g., metals and minerals), natural resources, processed materials, energy-related materials (e.g., fuels), construction materials, 3D printing materials, components and parts, agricultural products, biomaterials, recycled materials, and many others.

Energy resources module

[3504] In another example, the energy resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of energy resources. In embodiments, energy resources may include power grids, fossil fuel power plants, and the like.

Distributed energy resources module

[3505] In another example, the distributed energy resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of distributed energy resources. In embodiments, distributed energy resources may include solar panels, wind turbines, battery storage systems, and the like.

Renewable energy resources module

[3506] In another example, the renewable energy resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of renewable energy resources. In embodiments, renewable energy resources may include solar farms, wind farms, hydroelectric systems, and the like.

Health resources module

[3507] In another example, the health resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of health resources. In embodiments,

health resources may include diagnostic tools, patient records systems, telemedicine platforms, and the like.

Industrial resources module

[3508] In another example, the industrial resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of industrial resources. In embodiments, industrial resources may include assembly lines, heavy machinery, manufacturing systems, and the like.

Financial resources module

[3509] In another example, the financial resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of financial resources. In embodiments, financial resources may include trading platforms, risk management tools, credit scoring systems, and the like.

Transactional resources module

[3510] In another example, the transactional resources module may include a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of transactional resources. In embodiments, transactional resources may include payment gateways, point-of-sale systems, blockchain networks, and the like.

Transportation resources module

[3511] In another example, the transportation resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of transportation resources. In embodiments, transportation resources may include trains, buses, logistics platforms, autonomous vehicles, and the like.

Physical delivery resources module

[3512] In another example, the physical delivery resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of physical delivery resources. In embodiments, physical delivery resources may include delivery trucks, drones, parcel sorting systems, and the like.

Vehicular resources module

[3513] In another example, the vehicular resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of vehicular resources. In embodiments, vehicular resources may include electric vehicles, autonomous vehicles, taxis, forklifts, trucks, utility vehicles, material transport vehicles, automated guided vehicles (AGVs), refrigerated trucks, oil tanker trucks, emergency response vehicles, armored vehicles, and the like),

ships, tugboats, smart containers, trains, buses, aircraft, unmanned aerial vehicles (UAVS), electric bikes, electric scooters, and the like.

Shipping container resources module

[3514] In another example, the shipping container resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of shipping container resources. In embodiments, shipping container resources may include smart containers, container tracking systems, container storage systems, and the like.

Shipping resources module

[3515] In another example, the shipping resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of shipping resources. In embodiments, shipping resources may include cargo ships, freight tracking systems, shipping schedules, and the like.

Additive manufacturing resources module

[3516] In another example, the additive manufacturing resources module may implement a set of machine learning systems, a set of artificial intelligence systems, and/or a set of neural networks configured to optimize, by computing hardware, the provisioning of the set of additive manufacturing resources. In embodiments, additive manufacturing resources may include 3D printers, 3D printer subsystems, and the like.

Multi-factor resource optimization platform

[3517] At the resource layer 39000, AI capabilities for resource optimization enable a platform for providing multi-factor optimization of resources ("the multi-factor resource optimization platform"). Multi-factor resource optimization may refer to the process of achieving the best possible resource allocation outcome and/or solution that meets several objectives simultaneously. Multi-factor resource optimization objectives for operations of may include profitability, cost reduction, improved customer experience and satisfaction, employee satisfaction, sustainability, market expansion, innovation, operational efficiency, risk management, fraud prevention, safety, corporate social responsibility, and the like.

[3518] In some implementations, the resource layer 39000 may have AI capabilities for financial resource optimization associated with the use of AI workloads, including generative AI workloads in various applications. Generative AI workloads may refer to a set of tasks and/or processes that are handled or augmented by generative intelligence technologies. These workloads may comprise a range of activities where generative AI systems generate, enhance, or modify offerings and communications, content, experiences, and the like. Generative AI media workloads may also include personalization and/or localization workloads, such as when generative AI is used for tailoring content or experiences to individual or regional preferences, editing workloads, visual effects workloads, marketing generation workloads, or any other workloads described throughout this disclosure.

[3519] For example, generative AI may be utilized to generate maintenance instructions, reports, product designs, production plans, marketing materials, training materials, simulations, AR/VR experiences and the like; simulate process flows, vehicle routes, foot traffic, layouts, product demand scenarios, and the like; and many others. In embodiments, generative AI may be used to provide personalized and/or localized content, offers, promotions, experiences, and many others. In embodiments, financial costs for the use of generative AI workloads may include compute costs associated with generative AI, licensing costs associated with the use of generative AI (such as for the use of generative AI models trained on copyrighted data sets), and the like.

[3520] In embodiments, the compute costs may refer to the expenses associated with the computational resources needed to train and run AI models, including generative AI models. These costs may vary depending on model complexity, training data volume, training duration, hardware requirements (e.g., GPUs, TPUs, or the like), and cloud services/infrastructure, energy consumption, maintenance, scaling, software and/or development tools (e.g., such as software and development tools that facilitate generative AI model training and deployment), and the like.

[3521] In embodiments, the resource layer 39000 may have AI capabilities for computational resource optimization of generative AI workloads. In embodiments, the computational resources for generative AI media workloads may include hardware resources such as graphics processing units (GPUs), tensor processing units (TPUs), central processing units (CPUs), high-performance computing (HPC) clusters, and the like; software and frameworks such as machine learning frameworks, AI development platforms, data processing software, and the like; data storage solutions; networking infrastructure (e.g., high bandwidth networks); cloud computing platforms; and energy resources; among others.

[3522] In some implementations, the resource layer 39000 may have AI capabilities for computational resource optimization for workloads involving other types of computation, including, but not limited to HPC, quantum computing, distributed computing, gaming engines (e.g., the Unreal Engine), cloud computing, edge computing, other types of artificial intelligence models, neural networks, deep learning, and many others. In embodiments, these computational resources may be used in tandem with generative AI computational resources. For instance, generative AI models may be used with the Unreal Engine to generate casino gaming experiences and/or content.

[3523] Strategies for computational resource optimization (e.g., for advanced computation) at the resource layer 39000 may include the use of parallelization (e.g., data parallelization, model parallelization, pipeline parallelization, and hybrid parallelization), efficient algorithms (e.g., optimization algorithms, data sampling techniques, dimensionality reduction, quantization and pruning, sparse representations and the like), dynamic resource allocation based on computational load, optimization based on energy efficiency, load balancing, caching, and many others.

[3524] In embodiments, quality of service (QoS) optimization may converge with AI capabilities of resource layer 39000 to enable resource optimization for AI workloads. Performance optimization may ensure that the AI models meet speed and responsiveness requirements, which may involve model simplification, efficient algorithms, hardware acceleration (e.g., using GPUS

and TPUs), load balancing, and resource allocation. To achieve scalability, techniques such as elastic scaling may be leveraged to dynamically adjust computational resources based on demand. Failover mechanisms, redundancy, and error handling may be used to ensure the reliability and availability of the AI models. Latency management may be utilized to minimize delays in AI processing to meet real-time requirements. Techniques such as computational resource scheduling, job queue management, caching, and the like may be used to optimize the use of AI and other computational resources. Monitoring and analytics tools, AI model management techniques, and load testing and simulation may also be used to optimize the use of computational resources for AI workloads.

[3525] In some implementations, blockchain risk management may converge with AI capabilities 23002 of resource layer 39000 to enable resource optimization for generative AI workloads. Blockchain can facilitate the use of decentralized computational resources and decentralized computational resource marketplaces. Further, smart contracts can be used to automatically manage the allocation of computational resources based on predefined requirements and/or real-time needs. Computational resource usage and computational resource-related transactions may be recorded on a blockchain, providing a record of how resources are being used.

[3526] At the resource layer 39000, AI capabilities for resource optimization may utilize AI-driven predictive analytics, machine learning algorithms, AI-driven optimization algorithms, AI-driven digital twins that use simulations to test different scenarios or strategies, AI-driven autonomy, and any of the AI capabilities discussed throughout this disclosure.

[3527] In embodiments, the resource layer 39000 may manage and optimize a common set of resources that are used across different sets of enterprise and/or operational applications and/or use cases. This may include dynamic resource allocation and optimization strategies applicable to complex and varied environments such as industrial settings (e.g., manufacturing plants and warehouses), transportation infrastructures (e.g., airports and shipping ports), energy facilities (e.g., oil rigs and nuclear power plants), retail operations, datacenters, hospitality venues, and the like. By implementing a flexible resource management approach, the resource layer 39000 can dynamically adjust computational, logistical, and operational resources to meet the specific demands of these diverse environments.

[3528] The core functionality of the resource layer 39000 involves intelligent resource provisioning, which enables adaptability across different operational contexts. For instance, in a manufacturing environment, it might optimize machine utilization and energy consumption, while in a shipping context, it could manage container tracking, logistics routing, and real-time inventory management. Similarly, in datacenter and compute environments, the resource layer 39000 would focus on workload distribution, server resource allocation, and energy efficiency.

AI subsystem integrator system

[3529] Referring now to Fig. 280, an AI subsystem integrator system 39100 is illustrated in accordance with embodiments described herein. Artificial intelligence (AI) and machine learning (ML) systems disclosed herein may be used to facilitate the automation of data integration, networking integration and other technology integration by system 39100. For example, the AI

subsystem integrator system 39100 may train an integrator 39102 that includes at least one AI/ML system trained according to training workflows 39104 and training inputs 39106 to integrate a plurality of platforms and their respective subsystems A-n. For example, the training workflows 39104 and training inputs 39106 may utilize various methods and inputs disclosed herein and in the documents incorporated by reference (e.g., based on technical data sets, existing models, human supervision, deep learning on outcomes and/or robotic process automation learning on human execution of tasks, among others), across both the internal subsystems A-n within each of the platforms and across the different platforms and their subsystems. In the example provided, integrator 39102 includes a data integrator 39108, a network integrator 39110, an attribution integrator 39112, and other integrators 39114. For example, a set of AI systems may be trained to generate an intra-platform integration system 39116 that integrates various subsystems A-n within each of the platforms. Similarly, a set of AI systems may be trained with reference to the intra-platform integration system 39116 to generate a trans-platform integration system 39118 that facilitates interaction between the platforms. For example, integrator 39102 may learn the interfaces of a subsystem of a first platform and a subsystem of a second platform and to generate a data and network connection among them, as well as an interface by which a user, or an AI agent, may manage the data and network connection. Thus, disclosed herein is an artificial intelligence system that is configured and trained to generate a set of configurable integration capabilities across the internal subsystems and services of a set of platforms. Such an artificial intelligence system may be referred to herein as the "AI subsystem integrator system" and may integrate data and/or network communication, data storage (local and/or network), management interfaces, energy resources, and other capabilities. It may be noted that at scale an AI subsystem integrator system may generate a fully integrated or converged system that has massively parallel connectivity across the internal subsystems of the respective platforms, translating them from isolated platforms into a converged platform. This may evolve from various pairwise converged platforms among the various members of the platforms into a fully converged platform that encompasses subsystem elements of all of them.

[3530] In each case, a cross-service resource optimization system, which may comprise or employ a set of AI agents trained in ways noted above, may be used to manage, configure, deploy, provision, and/or optimize the respective subsystems or services that operate within a linked, integrated or converged system (the respective subsystems and/or services having previously been isolated to their respective standalone platform(s)).

[3531] For example, across any two or more subsystems or services across platforms may be optimized by measuring and allocating the energy that is used across the platforms, such as use of battery storage by devices, use of energy by GPUs in cloud computing or data centers (such as for generative AI workloads), and the like. Thus, in embodiments, methods and systems are provided for measuring and/or optimizing energy resources across the subsystems and services of at least two platforms.

[3532] As another example, across any two or more subsystems or services across platforms may be optimized by measuring and allocating the data storage resources that are used across the

platforms, such as use of storage by devices, by cloud computing or data centers (such as for generative AI workloads), and in particular by network resources, such as use of network storage for caching content to reduce latency and improve efficiency of use of network connectivity resources. Thus, in embodiments, methods and systems are provided for measuring and/or optimizing data storage resources across the subsystems and services of at least two platforms.

[3533] As another example, across any two or more subsystems or services across platforms may be optimized by measuring and allocating the networking resources that are used across the platforms, such as use of bandwidth of network resources by devices, by cloud computing or data centers, and the like to reduce latency and improve efficiency of use of the network or connectivity resources. Thus, in embodiments, methods and systems are provided for measuring and/or optimizing network resources across the subsystems and services of at least two platforms.

[3534] As another example, across any two or more subsystems or services across platforms may be optimized by measuring and allocating the computation storage resources that are used across the platforms, such as use of computation resources by devices, by cloud computing or data centers (such as for generative AI workloads), and the like. Thus, in embodiments, methods and systems are provided for measuring and/or optimizing computation resources across the subsystems and services of at least two platforms.

[3535] In embodiments, higher level management systems may collectively optimize across multiple factors, such as energy, computation and network resources, including use of artificial intelligence systems (alone or in combination with analytic models) that are trained as noted herein to achieve multifactor optimization across resource types.

Multiplatform attention management system

[3536] Referring now to Fig. 281, a higher-level multiplatform attention management system 39200 may be used to track and optimize attention across platforms (e.g., across integrated subsystems 39202, 39204, and 39206), such as by understanding where overall attention by customers (or by intelligent agents that operate on their behalf) may be increased by encouraging transitions across platforms, rather than maximizing only within a platform. This may include a multiplatform experience orchestration system 39208 for orchestrating experiences that are intentionally designed to encourage a user to move from platform to. The multiplatform attention management system 39200 may be used to monitor attention and provide input to a multiplatform attention orchestration system 39210 in order to improve the overall attention by users across the respective platforms. This may include any two platforms, or a converged platform among two or more platforms, or an overall converged platform in which the respective subsystems or services are fully integrated (including by AI as noted above).

[3537] In embodiments, a higher-level multiplatform revenue management system may be used to track and optimize revenue across platforms, such as by understanding where overall revenue by customers (or by intelligent agents that operate on their behalf) may be increased by encouraging transitions across platforms, rather than maximizing only within a platform. This may include the multiplatform experience orchestration system for orchestrating experiences (such as VR, entertainment, and other experiences) that are intentionally designed to encourage a user to move

from platform to platform. The revenue management system may be used to monitor attention and other factors that drive short-term revenue and long-term customer value and provide input to the orchestration system in order to improve the overall revenue from users across the respective platforms. This may include any two platforms, or a converged platform among two or more platforms, or an overall converged platform in which the respective subsystems or services are fully integrated (including by AI as noted above).

CLOSING THE LOOP BY BUILDING TRUST AND DISTRIBUTING DECISIONS IN DIGITAL TWINS TO ENABLE AGILITY AND EFFICIENCY

[3538] Throughout this disclosure and the documents incorporated by reference herein, various embodiments are provided of digital twin platforms and systems that leverage sensors and other information sources, robust connectivity, and intelligence systems to allow users to experience accurate representations of the states and activities of the many entities that are involved in the workflows of an individual, group or enterprise (e.g., a company, business unit, department, government, household, non-profit, or other enterprise). These include, among others, role-based digital twins that can be configured, or that self-configure, how data is collected, stored, processed and/or presented in ways that account for the roles of respective users (e.g., ones directed to financial users, strategic users, operational users, many others). Digital twins also include adaptive digital twins that leverage intelligence systems, such as artificial intelligence, analytic systems, expert systems, or hybrids involving various permutations and combinations thereof, to adapt how data is collected, stored, processed and/or presented based on context, such as based on the content of the information that is collected and processed. Among these adaptive digital twins, a subset, AI-driven digital twins, may use component artificial intelligence systems at any of the stages of the pipeline of information from a sensor or other basic information source to the presentation of content to a user. Artificial intelligence systems can provide outputs along a spectrum of autonomy, including: a) presenting reports, alerts, classifications, predictions, recommendations, analyses and other outputs for human review and action; b) undertaking human-supervised control of one or more aspects of a workflow, where the artificial intelligence system (acting as a co-pilot, assistant, intelligent agent, or the like to a human user) outputs a prediction and/or a recommendation for an action (which may be embodied in an instruction that can be processed by a system), and a human confirms, or adjusts, the nature of the action before it is implemented; and c) undertaking autonomous control of one or more systems or subsystems, where the artificial intelligence processes relevant information, completes necessary classifications, predictions and other decision-making steps (including where applicable, confirming risk management, governance, self-reflection or other steps) and triggers an action (such as by providing an instruction, control signal, or the like) to a system or component.

[3539] Along this spectrum from basic reporting to full autonomy, there is a varying extent to which there is a human being in the decision-making loop. Artificial intelligence systems have enormous promise to automate individual, group and enterprise workflows, providing more reliable, seamless monitoring of ongoing activities, faster response times to unexpected events, and more efficient execution of many types of tasks, among many other benefits. However, the extent

to which AI systems can be trusted can vary widely, such as based on: a) the stakes involved (e.g., are there risks to human life, health or property in making a wrong decision?); b) how the AI system was trained (e.g., are there reasons to believe that there is bias in the training data that will be carried forward into the outputs of the AI system); c) how the AI system is governed or supervised (e.g., is in compliance with regulatory, legal or other governance frameworks embedded in the workloads of the system as contemplated in this disclosure and the documents incorporated by reference herein, or is there a need for separate governance or supervision?); d) how the AI system performs, in absolute terms and relative to other available systems, including other AI systems, analytic systems, or human beings (experts, other individuals, groups, or crowds, for example); e) the quality, type and/or availability of data (which can include factoring in the cost of the data); f) the quality, type and/or availability of connectivity (which can include factoring in the cost of the connectivity); g) the quality, type and/or availability of computational resources (which can include factoring in the cost of the computational resources); h) the quality, type and/or availability of other necessary resources (e.g., energy); i) environmental or contextual parameters of the workflow to which the AI system will be applied (such as the complexity of a physical environment, the presence of human beings in proximity, or the like); and other factors. As an example among many, a well-trained AI system may be highly capable of driving a vehicle down a highway during daylight, doing so more safely than a human being who could be prone to sleepiness or distraction, but if computational, networking, or other resources are uncertain, it may be decided that the AI system should not be trusted to perform the task. Decisions about the extent to which one can trust an artificial intelligence system can be very difficult, and the extent of trust is a major factor in determining whether or not AI systems, and their many benefits, can be unlocked.

[3540] A major benefit of the various digital twin systems disclosed herein is that a digital twin platform, particularly one that is adaptive and self-configuring, can provide an ideal environment for decision making. Good decisions usually need a combination of a) fresh, accurate, relevant information; b) application of some degree of expertise; c) use of judgment to consider tradeoffs of risks; and/or d) leadership to cause implementation (which may encompass, among other things, taking on the consequences when outcomes are unfavorable). Each of these attributes can be supported by the data collection, storage and processing systems, connectivity systems, computational systems and intelligence systems described throughout this disclosure and the documents incorporated by reference herein. For example, a pipeline of IoT and edge devices and high QoS networks (adaptive networking) can pass granular, real-time sensor data about all entities in an operating environment (machines, components, humans, infrastructure, etc.), which manifests in visibility, such as for big data analytics and consulting use cases, as well as data for machine learning. Data architectures (e.g., adaptive sensing and data collection; edge query language; intelligent data storage and processing layers) and advanced computational architectures (hybrid cloud, edge, and quantum computing) can be used to optimize conditions for application of human and machine intelligence. For example, data architectures can adapt visual representations for human cognitive processing (including based on the skillset, role, experience, expertise, or cognitive parameters of a human, such measured by neurometrics, psychometrics, or

other systems) and/or prepare and stage data for artificial intelligence systems (including by cleansing, deduplication, generation of synthetic data, entity resolution, normalization and other processing, generation of various embeddings that use one AI system (e.g., a trained neural network) to transform and/or compressed input data into outputs that are suitable for efficient processing by AI systems, and the like). Computational architectures can be adapted to use an appropriate mix of cloud, edge and on-device processing systems (including various chipsets described herein, such as AI chipsets and hybrid chipsets involving AI and other functions integrated together). This can include quantum computing where beneficial. Systems integration, including various configurations of systems, systems-of-systems and the like into platform and infrastructure solutions (including PaaS and IaaS configurations), with various architectures, including service-oriented architectures, microservices architectures, and the like, can integrate all relevant systems across an individual's or group's set of activities, an enterprise, or an entire business ecosystem, such that a digital twin can have access to current state and activity parameters for many relevant entities. Data and sensor fusion, such as involving sensor data and many other data sets can assist in tracking or predicting outcomes (environmental data, market data, transaction data and many others). Advanced artificial intelligence techniques can combine with human insight to generate expert systems and models that understand and predict operational states, flows and more complex behaviors reflected in data (big data analytics and advanced artificial intelligence techniques), including compact models that are capable of operating effectively on sparse data and/or constrained computational architectures at the edge, as well as generative AI outputs that provide summaries, creative content, instructions, reports, and many other outputs from various inputs, including text, image, video, audio and multi-modal generative AI systems. AI-driven stakeholder interfaces can adaptively filter, prioritize, and distribute information for planning, simulation, decision making and operational action to appropriate stakeholders across an entire enterprise.

[3541] With the above stack of integrated systems, taking the form of the many embodiments disclosed herein and in the document incorporated herein by reference, distribution of control, of the various disparate systems, systems-of-systems, components, platforms, infrastructure elements, and other entities involved in the activities of an individual, group, or various levels and roles of enterprise can be enabled, with consideration being given, as noted above, to factors such as latency, security, safety, operational efficiency, reliability and trust.

[3542] The untapped power of the digital twin is its ability to enable collaboration, and to distribute decision making, to the right mix of human beings, artificial intelligence systems and other systems, whether in an enterprise or in the activities of an individual or group. The digital twin allows incremental change, experimentation, and easily reversible implementation, removing fear of closing the loop. For example, a digital twin can include a simulation environment that simulates, based on actual historical and current data, what outcomes most likely ensue from deployment of particular configurations of human and artificial intelligence elements in decision making and control loops. The digital twin can also provide a planning environment for deployment of artificial intelligence systems, as well as the data collection, data processing,

networking, connectivity, energy, and computational systems and architectures that support them; for example, an enterprise can, using the digital twin, plan various scenarios for obtaining resources to enable more powerful AI systems, comparing projected outcomes at various resource levels and mixes. The digital twin can enable graceful migration among humans, human-AI mixes (e.g., with agents and co-pilots) and autonomous AI systems, including deploying them in ways that can be rapidly adjusted, or reversed, such as based on changing performance capabilities, contextual and environmental factors, and events; for example, an AI system might be permitted to control transactions during normal daily operations, but it could be removed from the loop, or provided with greater human supervision, in the case of major shifts in an environment, such as during severe swings in the market, after an environmental catastrophe, or the like. Thus, provided herein is a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity. Also provided herein is a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations. The resources may include ones for data collection, data processing, networking, connectivity, energy, and computational systems and architectures that support them.

[3543] As noted throughout this disclosure and the documents incorporated by reference herein disclose the training of machine learning systems in various configurations, including deep learning, supervised learning, semi-supervised learning, and the like. In many such embodiments, a human expert provides input data that is used to train the AI system, such as by tagging data sets to assist in AI classification systems; by producing code, text, images, audio, video and other inputs that are used to train generative AI systems to predict a set of outputs from a prompt; by undertaking activities and behaviors that indicate preferences to train AI systems to generate recommendations; by configuring systems, platforms and architectures that can be used to train AI systems to generate similar configurations and recommendations; by undertaking decisions in various contexts and environments that are used to train AI systems to produce recommendations, make decisions and/or output control signals in similar contexts and environments, among others. In embodiments, any of the machine learning systems or other artificial intelligence systems disclosed herein or in the documents incorporated by reference herein may be embedded in a digital twin for the purpose of enabling training of the digital twin. This may include designating a set of users (such as domain experts, managers, operational supervisors, executives, or the like) as trainers for the AI system, including based on the role, core competency, expertise, experience, or other aspects of the set of users. In embodiments, this may include designating one set of users to provide input data for initial training of a particular AI system and another set of users (possibly overlapping in part with the first set) to supervise the outputs of the AI system. Thus, disclosed herein is a digital twin system with an interface system for designating a set of users as trainers for

the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. Also disclosed herein is a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin.

[3544] In embodiments, provided herein is a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows. This may be implemented at the level of an individual component or system, such as by allowing a user to configure, within the digital twin, the conditions under which a human (or group of humans, a combined system (human and AI co-pilot, for example), or an AI system alone will be designated as the authority for making a decision. This can include permanent settings ("never" or "always"), semi-permanent settings (e.g., ones that are reviewed on a systematic or episodic basis) and dynamic settings (i.e., where contextual conditions, environmental conditions, or other factors are processed in real time to determine what set of entities will be designed at the decision making authority). In embodiments, decision making authority may be distributed within the digital twin based on a broader decision making framework, such as a hierarchical framework (such as an enterprise hierarchy in which individual workers are organized into groups, with lines of reporting and supervision), a rules-based framework (such as a set of voting rules by which disputes about a decision are resolved), a collaborative framework (such as where prospective alternatives are presented and discussed by a set of contributors seeking a consensus recommendation, optionally involving simulations and planning capabilities noted above), a simulation framework (where, as noted above, the digital twin can simulate outcomes based on historical and real-time data), an enterprise planning framework (which may include integration of the digital twin with enterprise planning systems and dashboard), a competitive framework (such as where alternative options are configured to compete with each other to determine which is superior, optionally involving genetic programming or other evolutionary computing systems, as noted elsewhere in this disclosure), a peer-to-peer framework (e.g., where decisions are negotiated bi-laterally or multi-laterally among entities involved); an algorithmic framework (such as where decisions are reached by a defined set of stages, possibly involving various other frameworks as hybrids in parallel, in series, in iterative loops, or the like), a principles-based framework (such as where a set of personal, group, or enterprise principles (e.g., ethical principles, values, mission statements, or the like) are codified for use when certain contextual decisions are presented), or others. As an example, many daily operational decisions might be configured to be executed autonomously by AI systems, but if those decisions are determined to have an aggregate impact on an element of the mission of an enterprise (such as a mission to continue to develop and employ a set of individuals), then the digital twin can highlight the need to bring an appropriate set of human decision makers into the loop. Thus, in embodiments, provided herein is a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-

based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework.

[3545] In embodiments, a digital twin may display various metrics to facilitate decisions about what sets and configurations of human and AI systems should be selected for a given system, system-of-systems, component, or the like. These metrics may be calculated using the data available to the digital twin through its processing functions (such as real time information about the user, available human workers, resource availability for AI systems, contextual and environmental information, information about states and activities of the entities involved in various workflows and the like) as well as by using external data sources, such as data indicating performance metrics for individuals (which may include the general population or domain experts, groups (which may include expert groups, crowds (including by crowdsourcing as disclosed elsewhere herein)) or for AI systems (including how the systems perform under various resource conditions). Metrics may include indicators of training history, such as educational experience, work experience, and job reviews, among many others, for humans. Metrics may include training data or metadata for artificial intelligence systems, including the size of the training data set, the vintage of the training data set, the configuration of neural networks used in training, the presence or absence of synthetic data in the training data set, metrics indicating the quality of the training data set (including indicators of bias, autocorrelation, heteroskedasticity or other statistical or econometric indicators), metrics about the expertise or capabilities of the human beings used to seed and/or supervise the AI systems, and many others. As an example, two similar generative AI systems may be presented along with indicators of the time period over which they were trained, allowing a user to evaluate whether the AI systems are likely to miss important context (e.g., where training data from a time period outside one of the sets is likely to have a major influence on the outcome) or generate hallucinations (such as where a spike in unusual training data during a training period may have unduly swayed results). Thus, in embodiments, provided herein is a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. Also provided herein, in embodiments, is a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. Also provided herein, in embodiments, is a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin.

[3546] In embodiments, metrics provided within the digital twin may include performance metrics for one or more sets of human and AI entities. Human performance metrics may include metrics for individuals, experts, and groups (including crowds) of the many types, and collected by the many methods and systems, disclosed herein and in the documents incorporated by reference. For individuals, these may include various cognitive, neurometric, psychometric, emotional, attention and other metrics, including ones disclosed herein and in the documents incorporated herein, including collected by neuro-metric measurement systems (e.g., EEG, fMRI and other scanning

systems), genomic and transportomics systems, psychometric testing, physiological monitoring systems (including wearable sensors, cameras, IoT systems and many others as disclosed throughout this disclosure), systems for detecting emotional state (e.g., anxiety, distress, anger, calm and others, and others), systems for measuring attention, and others. Human performance metrics may also include metrics of task performance, including quality metrics, output metrics, job assessment metrics, indicators of specific skills, expertise or competencies (e.g., educational credentials, publications and certifications), indicators of success (e.g., rates of return of a business led by the individual) and many others. Human metrics can include metrics for groups, including crowds; for example, outcome metrics from a crowdsourcing system that accumulates the "wisdom of the crowd" for a set of predictions, ideas, solutions, or the like can be compared to those of individual experts and those of various AI systems (including standalone systems and human-AI combinations). AI performance metrics may include a wide range of metrics, including metrics of predictive accuracy, metrics indicating outcomes from use of the systems (including one's accounting for context of use), metrics of computational speed and latency, metrics of resource utilization (e.g., computation, energy, network resources, and the like), and many others. In embodiments, the digital twin may present comparative metrics across human, human-AI combinations and autonomous AI systems (including multiple options for each), so that a user of the digital twin can select, accounting for context of use, an appropriate configuration. Thus, provided herein, in embodiments, is a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. Also provided herein, in embodiments, is a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. Also provided herein, in embodiments, is a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. Also provided herein, in embodiments, is a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin.

[3547] It should be noted that an intelligent agent (referred to in some cases herein as an opportunity miner) can be used to discover sets of available sets of humans, human-AI combinations, and/or artificial intelligence systems that may be capable of improving outcomes of one or more systems, systems-of-systems or other entities represented in the digital twin. In embodiments, a user of a digital twin may configure a set of intelligent agents to undertake discovery based on prioritization by the user, such as where a user flags elements of an enterprise workflow that are perceived (or determined by metrics) to be most in need of improvement, or most likely to benefit from artificial intelligence. Thus, a digital twin system may be used as a tool for exploration of applications of artificial intelligence as well as for configuring decisions around deployment of artificial intelligence once discovered. Thus, provided herein is an intelligent agent

that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. Also provided herein is a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function.

[3548] In embodiments, a user of a digital twin may configure a set of intelligent agents to undertake discovery based on prioritization by the user, such as where a user flags elements of an enterprise workflow that are perceived (or determined by metrics) to be most in need of improvement, or most likely to benefit from artificial intelligence.

[3549] In embodiments, an artificial intelligence system may be trained, such as based on a set of human decisions, a set of outcomes, or the like, and used to configure (or recommend configuration of) deployment of one or more appropriate sets of human, human-AI combinations and autonomous AI systems. This deployment configuration system for artificial intelligence may be integrated or embedded as an enabling service or utility within a digital twin, or it may be a standalone system used to determine what options are to be presented to a user of a digital twin. Thus, provided herein is an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. Also provided herein is a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems.

[3550]

[3551] In embodiments, the various discovery and configuration systems noted above may be based on general metrics of performance (e.g., comparing how humans in general perform relative to artificial intelligence systems at performing particular tasks). In other embodiments, the comparison may be highly contextual and situational, such as comparing the capabilities of specific sets of individuals that will be available at the time of decision making and the capabilities of artificial intelligence systems that will be available, either to operate as standalone systems or to work in human-AI combinations (including the general performance capabilities of the AI systems, and also the available computational, networking, energy and other resources that will be needed to run them). Situational comparison can include contextual factors, such as time of day, workforce availability, market and environmental data, and many others, such that a discovery or configuration system can recommend an appropriate configuration for a particular system, at a given place and time. For example, a moderately expensive AI system that performs well autonomously may be recommended or selected for an overnight control task, where human experts are not likely to be available and where computational, energy, or other resources are less constrained (and less expensive). An intelligent agent, as noted above, may be configured to generate a recommendation, or a configuration, based on multi-factor optimization, using various

decision frameworks noted above, including being trained on a set of human decisions and/or on feedback of outcomes from past configuration actions.

[3552] Thus, provided herein is an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed.

[3553] In embodiments, as noted throughout this disclosure, smart contracts can be integrated with one or more systems of a digital twin, such that they operate on input data (e.g., detection of various events) and produce outputs that reflect transactional terms embedded in the smart contract. This may include, for example, a set of smart contracts that set liability terms relating to the consequences of the outputs of a system to which the smart contract relates. For example, a smart contract may be configured such that the provider of an AI system embedded in the digital twin (or in a system to which the digital twin has access) agrees to provide indemnification, proof of insurance, acceptance of liability, or the like for some set of consequences of using the AI system. Such terms can also include various limitations (e.g., constraints on permitted use; limitations of liability), exclusions, restrictions, and the like, thus providing, in the digital twin, a mechanism for expressly allocating the consequences of selection of a particular mix of human and AI systems to the appropriate person or organization. By referencing the performance metrics noted above, operating in real time on operational and other data, a decision maker can be aided to make a rational decision about implementation of AI, guided by a more reliable estimation of the expected value of implementing a particular mix of human and AI systems, rather than being governed by emotional factors. Thus, provided herein are methods and systems having a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin. The system displayed in the digital twin may be an artificial intelligence system.

[3554] As good decisions made within a digital twin environment or other decision making platform lead to better outcomes, trust in the reliability of machine recommendations increases, which can be tested and validated by outcome metrics. Individuals, groups and enterprises can gradually migrate decision making from centralized authority to the operating environment and from human action through various degrees of supervision to autonomy, with appropriate checks and balances at all stages (automated governance), reversing direction when the situation calls for it. Over time as trust builds in the capabilities of artificial intelligence, or human-AI combinations, to outperform humans alone, an enterprise can increasingly close loops among information technology, operations technology, and artificial intelligence technologies, enabling low-latency, highly efficient, well governed autonomous response to changing conditions at the operational

level. Operators that don't ultimately close the loop will be slower than their competitors to respond to changing conditions, less efficient in their operations, and less effective in their decision making.

[3555] Individuals, groups and enterprises can unlock great benefits by properly distributing decision making and control where and when it is most effective throughout an enterprise, enabled by the integrated stack of technologies that enable a digital twin, each enhanced by artificial intelligence as described throughout this disclosure and the documents incorporated by reference herein. The true power of the convergence of information technology and operations technology (IT/OT convergence) can be realized when decision making is optimally distributed across humans and machines, an outcome that is achieved by progressively building trust in intelligence technologies and the people who supervise them. Organizations that progress more rapidly to a more optimal state will have significant competitive advantages in operational efficiency (particularly through closed loop automation) and agility to respond to shifting market and competitive dynamics.

[3556] In embodiments, various embodiments of digital twins, intelligent agents for discovery and configuration, decision making frameworks and other methods and systems disclosed herein may be used to facilitate configuration and allocation of decision making across the various resources, entities, activities, operations, transactions, offerings and workflows involved in various manufacturing, energy production, or other industrial environments that are described throughout this disclosure and the documents incorporated by reference herein.

[3557] In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for

the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based

platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based

platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin having a simulation environment that simulates, based on data for the entities, states and activities processed and displayed by the digital twin, a simulation of a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in the decision making or control workflows of an entity and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3558] In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the

artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at

least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin that includes a

planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability

factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin that includes a planning system for deployment of a set of artificial intelligence systems and the resources required to enable the artificial intelligence systems, wherein the planning system provides a prediction of the outcomes of deployment of at least a plurality of distinctive sets of artificial intelligence systems or a plurality of distinctive sets of resource configurations and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3559] In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided

herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone

artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as trainers for the creation of an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3560] In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform

having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a digital twin system having an integrated intelligent

agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system with an interface system for designating a set of users as supervisors for an artificial intelligence system that is represented in the digital twin, wherein the artificial intelligence system operates on the data that is used to populate the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3561] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision

making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of

entities and workflows and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of options for distribution of decision making authority or control across various systems, systems-of-systems or components across a set of entities and workflows and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3562] In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In

embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin.

In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a

principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin wherein authority for decision making about what set of entities among human beings, human-AI systems, or autonomous AI systems is distributed within the digital twin based on a decision making framework selected from a hierarchical framework, a rules-based framework, a simulation framework, an enterprise planning framework, an algorithmic framework, a principles-based framework, a collaborative framework, a peer-to-peer framework, or a competitive framework and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3563] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays

a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment

parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3564] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays

a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having

a digital twin system that displays a set of training metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3565] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the

digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of training metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3566] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin

system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of humans that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that

sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3567] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital

twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of hybrid human-AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3568] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically

configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of performance metrics for a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3569] In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial

intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system that displays a set of comparative performance metrics among one or more of a set of individual humans, a human-led enterprise, a crowd of humans, a human-AI combination system, or a set of AI systems that are available to perform a set of decision making tasks for a set of entities presented in the digital twin and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3570] In embodiments, provided herein is an AI-based platform having an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI

systems and standalone artificial intelligence systems that are capable of performing a desired function and a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having an intelligent agent that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3571] In embodiments, provided herein is an AI-based platform having a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. In embodiments, provided herein is an AI-based platform having a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired

function. and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system having an embedded intelligent agent system that automatically discovers sets of available systems, among human systems, combined human-AI systems and standalone artificial intelligence systems that are capable of performing a desired function. and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3572] In embodiments, provided herein is an AI-based platform having an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems and a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having an intelligent agent that operates as

a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3573] In embodiments, provided herein is an AI-based platform having a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems. In embodiments, provided herein is an AI-based platform having a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, and an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having a digital twin system having an integrated intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3574] In embodiments, provided herein is an AI-based platform having an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed. In embodiments, provided herein is an AI-based platform having an intelligent agent that operates as a deployment configuration system to at least one of recommend a set of deployment configuration parameters for or automatically configure a set of deployment parameters for a set of human, human-AI combinations or autonomous AI systems, wherein a deployment configuration is determined at least in part on a set of situational factors, the situational factors being among one or more of a set of availability factors for human resources or artificial intelligence resources, a set of capability factors for human resources or artificial intelligence resources, a set of resource availability factors for enabling a set of human resources or artificial intelligence resources, or a set of contextual or environmental factors for the system to which the configuration is to be deployed and a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

[3575] In embodiments, provided herein is an AI-based platform having a smart contract system embedded in a digital twin that sets terms and conditions for liability resulting from the output of a system that is displayed in the digital twin.

VISUAL SEARCH WITHIN DIGITAL TWINS

[3576] In embodiments, a visual search engine may be embedded within a digital twin, wherein a user may enter an image-based query (such as by pasting an image, a set of image metadata, or the like, by drawing or otherwise creating the image, or by describing the image with a set of keywords), and the visual search engine will process the query and a set of data structures supporting the digital twin to find a set of potentially matching elements in the digital twin (such as by finding matching image identifying information, matching tags, matching or correlated metadata, or the like, or by finding images having characteristics (such as shape, color, semantic or symbolic meaning, or other attributes) that match or are correlated to the query. In embodiments, matching elements may be listed (such as in a menu bar) and/or highlighted in the visual representation of the digital twin, such as by changes in font, color, brightness, or the like. In one non-limiting example, a user of a digital twin of a building, such as an architect, designer, builder, owner, operator, tenant, visitor, or other users, may enter a picture of a space that is believed to be within the building, containing elements such as fixtures, furnishings, equipment and the like, and the visual search engine may parse the picture, find matching elements in the digital twin of the building, and present one or more options that represent likely location of the space, along with

other data or representations of the space, including guidance how to reach the space, additional data on contents, and the like. For example, a visiting maintenance professional may enter a picture of a machine that was captured on an earlier visit by a colleague or by a customer, and the visual search engine may find where the machine is now located in the digital twin and route the professional to the current location. In embodiments, the visual search engine may use artificial intelligence, in one of the many forms described herein or in the documents incorporated by reference, such as to parse queries (in natural language or in image formats); to recognize and classify objects, faces, spaces, and other visual elements, to match or correlate images and image content, text, or the like; to rank results or recommendations, such as based on the likelihood of relevance, including based on past user responses, ratings, utilization, or other measures of feedback; to optimize the presentation of results; and many other elements.

[3577] In embodiments, a digital twin of an environment, such as the interior of a home, may be automatically constructed from an arbitrary set of photos that are taken over time, such as snapshots of people, pets, and other subjects contained in the photo library of a consumer, whereby common elements within the photographs are linked by a set of artificial intelligence systems, such as object recognition systems, subsequent to which the comment elements enable a user to orient, resize, and otherwise manipulate image elements to provide a stitched view of the environment. In embodiments, a large training data set may be captured as expert users stitch view of multiple environments, such that a robotic process automation system may be trained to automatically stitch a digital twin of an environment upon receiving an input data set. In embodiments, incremental information may be added to the digital twin, such as additional images as they are captured and other content, such as information from other data sources about the environment, including data about or from IoT devices in the environment, about objects purchased for the environment, about workflows in the environment, about individuals in the environment (such as user behavior data) and the like. In embodiments, image information and other information about the environment may be captured in a blockchain. The blockchain may support a distributed ledger of transactions or events related to the environment. The blockchain may support a smart contract as described elsewhere in this disclosure. The blockchain may be augmented by a set of interfaces, such as for controlling access to the information, such as to enable private portions and public portions. In embodiments, public portions may allow third parties, such as vendors, suppliers, service providers, technology providers and others to obtain sufficient relevant information from the digital twin to meet specific needs of an owner or operator of the environment, or a worker within it, without providing access to the entire set of information contained in the digital twin.

INDUSTRIAL TECHNOLOGY STACK

[3578] FIG. 282 is a diagram of an industrial technology stack in accordance with some example embodiments. As shown, the industrial technology stack includes a set of layers, including a governance layer, an enterprise layer, an offering layer, a transaction layer, an operations layer, a network layer, a data layer, and a resource layer. The industrial technology stack may be used in an industrial internet-of-things (IIoT) environment, for example, to enable devices in an industrial environment (*e.g.*, a manufacturing facility) to cooperatively discover and assign resources,

orchestrate the application of computational resources to industrial resources and uses, to monitor and manage industrial processes, and/or to maintain the industrial environment (*e.g.*, diagnosing and detecting issues with the machinery of the industrial environment and conducting maintenance and repairs).

[3579] In some embodiments, the industrial technology stack of the IIoT environment may include and/or resemble the AI convergence system of systems discussed herein and shown in FIGS. 271-281. Alternatively or additionally, the industrial technology stack of the IIoT environment may include components that are not included in the AI convergence system of systems discussed herein. The industrial technology stack of the IIoT environment may include components that do not include artificial intelligence, such as traditional software elements, hardware elements and machinery that do not include software of any kind, and/or other components such as chemical compositions and/or biological processes

[3580] In embodiments, the industrial technology stack of the IIoT environment includes a governance layer. The governance layer of the industrial technology stack provides governance of other portions of the industrial technology stack and/or the industrial environment. For example, the governance layer may monitor the resources of the industrial environment (*e.g.*, manufacturing machines, robots, computers, stores of supplies, sensors, data stores, communication devices, and the like) to ensure that governance policies are being met. In example embodiments, the governance policies include at least one of: conserving the consumption of resources, improving the efficiency of industrial processes, promoting productivity of employees of the industrial environment, reducing greenhouse gas emissions, and/or ensuring compliance with at least one law, regulation, moral principle, ethical obligation, organizational objective, or the like. The governance layer may use technologies such as artificial intelligence models for embedded policies and governance; policy automation; safety and compliance automation; and reporting automation to achieve automated governance of the industrial operations and the industrial environment as a whole.

[3581] In embodiments, the industrial technology stack of the IIoT environment includes an enterprise layer. The enterprise layer of the industrial technology stack provides enterprise-level management of the industrial environment. For example, the enterprise-level management may include holistically allocating resources of an industrial enterprise (*e.g.*, operating space, storage space, energy, supplies, processing time for computing and industrial machinery, funds, human attention, or the like) to the industrial processes and equipment of the industrial environment. The enterprise-level management may select, analyze, and/or adapt an organizational structure of the enterprise and/or industrial environment to various entities within the enterprise (*e.g.*, departments, branches, projects, executives, groups of workers, or the like). The enterprise layer may use technologies such as artificial intelligence models for contextual simulation and forecasting, executive digital twins, IIoT platforms, big data and analytics, and advanced operational control over the industrial environment and/or an industrial metaverse.

[3582] In embodiments, the industrial technology stack of the IIoT environment includes an offering layer. The offering layer of the industrial technology stack provides for the formulation,

evaluation, and selection of system offerings by the industrial environment. The offerings may include, for example, contracts for industrial processes; the acquisition, sale, exchange, and/or processing of industrial products, equipment, and/or supplies; research or manufacturing projects related to the operations and/or interests of the industrial environment; and/or the formation of coalitions of industry partners for various objectives, such as establishing technology or industry standards. The processing of the offerings may include, for example, a selection or description of external contracts or opportunities for participation by the industrial enterprise; a formulation, analysis, and/or release of external offerings by the industrial enterprise for the participation of other enterprises; and/or the selection, description, formulation, analysis, and/or release of internal offerings between or among various sets of entities within the industrial enterprise. The offerings layer may use technologies such as expert systems and generative AI, industrial metaverse, systems integration, industrial control systems, and industrial automation of an industrial service ecosystem.

[3583] In embodiments, the industrial technology stack of the IIoT environment includes a transaction layer. The transactions layer may formulate, negotiate, execute, fulfill, and/or report transactions related to the offerings of the offering layer. For example, the transactions may include a purchase or sale of one or more assets, wherein the purchase or sale includes a transfer of value in exchange for the one or more assets; a mortgage, lease, or loan of resources; and/or a conveyance, transport, and/or delivery of a physical or digital asset. The transaction layer use techniques such as smart contracts for the execution of transactions for the industrial enterprise. The transaction layer may use technologies such as artificial intelligence models for discovery, generation, and optimization of transactions; smart contracts and distributed ledgers; a transaction-related system-of-systems; automated transaction orchestration; and/or an automated industrial marketplace.

[3584] In embodiments, the industrial technology stack of the IIoT environment includes an operations layer. The operations layer provides components for conducting a set of operations related to the other layers of the industrial technology stack and the industrial environment. The operations may include, for example, executing the steps of an industrial process such as a manufacturing, refining, materials processing, inspection, assembly, or delivery process; monitoring, assessing, diagnosing, maintaining, repairing, and/or cleaning various machinery and/or equipment of the industrial environment; and/or generating, organizing, interconnecting, and/or maintaining one or more components of the industrial technology stack. The operations layer may use technologies such as artificial intelligence models for routing, control, optimization, and generation; automated monitoring and maintenance platforms; asset and infrastructure management; robotics and process automation; and converged, AI-based IT/OT workflow orchestration.

[3585] In embodiments, the industrial technology stack of the IIoT environment includes a network layer. The network layer provides connectivity and communicating among various components of the industrial environment, such as industrial machines, computers such as servers and workstations, sensors, personal communication devices such as mobile phones, and software

components such as artificial intelligence models and conventional algorithms of the industrial technology stack. The network layer may include a diverse and dynamic set of communication technologies, such as wired networking (*e.g.*, Ethernet or fiber-optic connections), large-scale wireless communication (*e.g.*, Wi-Fi and/or cellular), small-scale wireless communication (*e.g.*, Bluetooth and/or infrared), and satellite communication. The network layer may connect one or more resources of the industrial environment to one or more external communication partners, such as the Internet, and/or may isolate one or more resources of the industrial environment from communicating with external communication partners. The network layer may use technologies such as artificial intelligence models for adaptive networking; edge and/or cloud communication; cellular, Wi-Fi, ORAN, and/or Bluetooth communication protocols; Industrial Internet-of-Things (IIoT)-related technologies, such as IIoT communication protocols; and/or intelligence edge, advanced networking, and/or IIoT-in-a-Box architectures.

[3586] In embodiments, the industrial technology stack of the IIoT environment includes a data layer. The data layer provides and manages the data storage capabilities of the industrial enterprise, such as storing, curating, indexing, searching, querying, structuring, providing, refreshing, compiling, transmitting, altering, monitoring, and/or deleting data. The data may include, for example, files in a file system, records of one or more databases, data objects in formats such as XML or JSON, executable binaries, scripts, portions of memory, portions of one or more caches, or the like. In some example embodiments, the data layer performs tasks such as generating data stores; preparing or adapting data stores for various industrial processes; filling data stores with data; and/or restructuring sored data for industrial uses and/or reports. The data layer may use technologies such as artificial intelligence models for sensor fusion, sensor and wearable devices, mobile data collectors, APIs, SOA and distributed data, and context-aware industrial IoT and sensor fusion.

[3587] In embodiments, the industrial technology stack of the IIoT environment includes a resource layer. The resource layer may acquire, manage, monitor, control, transport, and/or maintain one or more specific resources of the industrial environment, such as a piece of manufacturing equipment, a robot, a supply of a material such as fuel, or an industrial product such as a synthesized composition or an assembled device. The resource layer may use the resource of the industrial technology stack to perform tasks such as monitoring a status and/or condition of the resource; monitoring a status and/or condition of an industrial process involving the resources, such as an industrial workflow; arranging and/or scheduling maintenance, repair, and/or resupply of the resource; and/or generating reports of a status or condition of the resource or industrial processes related thereto. The resource layer may use technologies such as artificial intelligence models for resource optimization, advanced computation, switches and multiplexers, vibration and other sensor data collection, and machine-level data collection and resource optimization.

[3588] FIG. 283 is a diagram of a set of examples of industrial modules of the industrial technology stack. The industrial technology stack may include one or more industrial modules that perform and/or adapt functions of the industrial technology stack for a particular industrial environment, such as a particular manufacturing facility, a particular manufacturing process, and/or a particular

manufacturing marketplace. The industrial modules may include and/or perform at least one of: automated governance of industrial operations; AI-based advanced operational control and AI-driven industrial metaverse; industrial automation and industrial service ecosystem service; a provision and/or participation in an automated industrial marketplace; converged, AI-based IT/OT workflow orchestration; intelligent edge, advanced networking, and industrial IoT-in-a-box architectures; context-aware industrial IoT and sensor fusion; and/or machine-level data collection and resource optimization.

HARDWARE, SOFTWARE, AND GENERAL-PURPOSE SYSTEMS

[3589] Special-purpose systems include hardware and/or software and may be described in terms of an apparatus, a method, or a computer-readable medium. In various embodiments, functionality may be apportioned differently between software and hardware. For example, some functionality may be implemented by hardware in one embodiment and by software in another embodiment. Further, software may be encoded by hardware structures, and hardware may be defined by software, such as in software-defined networking or software-defined radio.

[3590] The methods and/or processes described in the disclosure, and steps associated therewith, may be realized in hardware, software or any combination of hardware and software suitable for a particular application. The hardware may include a general-purpose computer and/or dedicated computing device or specific computing device or particular aspect or component of a specific computing device. The processes may be realized in one or more microprocessors, microcontrollers, embedded microcontrollers, programmable digital signal processors or other programmable devices, along with internal and/or external memory. The processes may also, or instead, be embodied in an application specific integrated circuit, a programmable gate array, programmable array logic, or any other device or combination of devices that may be configured to process electronic signals. It will further be appreciated that one or more of the processes may be realized as a computer executable code capable of being executed on a machine-readable medium.

[3591] In this application, including the claims, the term module refers to a special-purpose system. The module may be implemented by one or more special-purpose systems. The one or more special-purpose systems may also implement some or all of the other modules. In this application, including the claims, the term module may be replaced with the terms "controller" or "circuit." In this application, including the claims, the term platform refers to one or more modules that offer a set of functions. In this application, including the claims, the term system may be used interchangeably with module or with the term special-purpose system.

[3592] The special-purpose system may be directed or controlled by an operator. The special-purpose system may be hosted by one or more of assets owned by the operator, assets leased by the operator, and third-party assets. The assets may be referred to as a private, community, or hybrid cloud computing network or cloud computing environment. For example, the special-purpose system may be partially or fully hosted by a third party offering software as a service (SaaS), platform as a service (PaaS), and/or infrastructure as a service (IaaS). The special-purpose system may be implemented using agile development and operations (DevOps) principles. In

example embodiments, some or all of the special-purpose system may be implemented in a multiple-environment architecture. For example, the multiple environments may include one or more production environments, one or more integration environments, one or more development environments, etc.

[3593] A special-purpose system may be partially or fully implemented using or by a mobile device. Examples of mobile devices include navigation devices, cell phones, smart phones, mobile phones, mobile personal digital assistants, palmtops, netbooks, pagers, electronic book readers, tablets, music players, etc. A special-purpose system may be partially or fully implemented using or by a network device. Examples of network devices include switches, routers, firewalls, gateways, hubs, base stations, access points, repeaters, head-ends, user equipment, cell sites, antennas, towers, etc.

[3594] A special-purpose system may be partially or fully implemented using a computer having a variety of form factors and other characteristics. For example, the computer may be characterized as a personal computer, as a server, etc. The computer may be portable, as in the case of a laptop, netbook, etc. The computer may or may not have any output device, such as a monitor, line printer, liquid crystal display (LCD), light emitting diodes (LEDs), etc. The computer may or may not have any input device, such as a keyboard, mouse, touchpad, trackpad, computer vision system, barcode scanner, button array, etc. The computer may run a general-purpose operating system, such as the WINDOWS operating system from Microsoft Corporation, the MACOS operating system from Apple, Inc., or a variant of the LINUX operating system. Examples of servers include a file server, print server, domain server, internet server, intranet server, cloud server, infrastructure-as-a-service server, platform-as-a-service server, web server, secondary server, host server, distributed server, failover server, and backup server.

[3595] The term hardware encompasses components such as processing hardware, storage hardware, networking hardware, and other general-purpose and special-purpose components. Note that these are not mutually-exclusive categories. For example, processing hardware may integrate storage hardware and vice versa.

[3596] Examples of a component are integrated circuits (ICs), application specific integrated circuit (ASICs), digital circuit elements, analog circuit elements, combinational logic circuits, gate arrays such as field programmable gate arrays (FPGAs), digital signal processors (DSPs), complex programmable logic devices (CPLDs), etc.

[3597] Multiple components of the hardware may be integrated, such as on a single die, in a single package, or on a single printed circuit board or logic board. For example, multiple components of the hardware may be implemented as a system-on-chip. A component, or a set of integrated components, may be referred to as a chip, chipset, chiplet, or chip stack. Examples of a system-on-chip include a radio frequency (RF) system-on-chip, an artificial intelligence (AI) system-on-chip, a video processing system-on-chip, an organ-on-chip, a quantum algorithm system-on-chip, etc.

[3598] The hardware may integrate and/or receive signals from sensors. The sensors may allow observation and measurement of conditions including temperature, pressure, wear, light, humidity, deformation, expansion, contraction, deflection, bending, stress, strain, load-bearing, shrinkage,

power, energy, mass, location, temperature, humidity, pressure, viscosity, liquid flow, chemical/gas presence, sound, and air quality. A sensor may include image and/or video capture in visible and/or non-visible (such as thermal) wavelengths, such as a charge-coupled device (CCD) or complementary metal-oxide semiconductor (CMOS) sensor.

[3599] The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software, program codes, and/or instructions on a processor. The present disclosure may be implemented as a method on the machine, as a system or apparatus as part of or in relation to the machine, or as a computer program product embodied in a computer readable medium executing on one or more of the machines. In example embodiments, the processor may be part of a server, cloud server, client, network infrastructure, mobile computing platform, stationary computing platform, or other computing platforms. A processor may be any kind of computational or processing device capable of executing program instructions, codes, binary instructions and the like, including a central processing unit (CPU), a general processing unit (GPU), a logic board, a chip (e.g., a graphics chip, a video processing chip, a data compression chip, or the like), a chipset, a controller, a system-on-chip (e.g., an RF system on chip, an AI system on chip, a video processing system on chip, or others), an integrated circuit, an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), an approximate computing processor, a quantum computing processor, a parallel computing processor, a neural network processor, an approximate computing processor, a quantum computing processor, a parallel computing processor, a neural network processor, a signal processor, a digital processor, a data processor, an embedded processor, a microprocessor, and a co-processor. The co-processor may provide additional processing functions and/or optimizations, such as for speed or power consumption, or other type of processor. Examples of co-processors include a math co-processor, a graphics co-processor, a communication co-processor, a video co-processor, and an artificial intelligence (AI) co-processor. The processor may be or may include a signal processor, digital processor, data processor, embedded processor, microprocessor or any variant such as a co-processor (math co-processor, graphic co-processor, communication co-processor, video co-processor, AI co-processor, and the like) and the like that may directly or indirectly facilitate execution of program code or program instructions stored thereon. In addition, the processor may enable execution of multiple programs, threads, and codes. The threads may be executed simultaneously to enhance the performance of the processor and to facilitate simultaneous operations of the application. By way of implementation, methods, program codes, program instructions and the like described herein may be implemented in one or more threads. The thread may spawn other threads that may have assigned priorities associated with them; the processor may execute these threads based on priority or any other order based on instructions provided in the program code.

[3600] The processor, or any machine utilizing one, may include non-transitory memory that stores methods, codes, instructions and programs as described herein and elsewhere. The processor may access a non-transitory storage medium through an interface that may store methods, codes, and instructions as described herein and elsewhere. The storage medium associated with the processor

for storing methods, programs, codes, program instructions or other type of instructions capable of being executed by the computing or processing device may include but may not be limited to one or more of a CD-ROM, DVD, memory, hard disk, flash drive, RAM, ROM, cache, network-attached storage, server-based storage, and the like. The computer software, program codes, and/or instructions may be stored and/or accessed on machine readable media that may include: computer components, devices, and recording media that retain digital data used for computing for some interval of time; semiconductor storage known as random access memory (RAM); mass storage typically for more permanent storage, such as optical discs, forms of magnetic storage like hard disks, tapes, drums, cards and other types; processor registers, cache memory, volatile memory, non-volatile memory; optical storage such as CD, DVD; removable media such as flash memory (e.g., USB sticks or keys), floppy disks, magnetic tape, paper tape, punch cards, standalone RAM disks, Zip drives, removable mass storage, off-line, and the like; other computer memory such as dynamic memory, static memory, read/write storage, mutable storage, read only, random access, sequential access, location addressable, file addressable, content addressable, network attached storage, storage area network, bar codes, magnetic ink, network-attached storage, network storage, NVME-accessible storage, PCIE connected storage, distributed storage, and the like.

[3601] A processor may include one or more cores that may enhance speed and performance of a multiprocessor. In example embodiments, the process may be a dual core processor, quad core processors, other chip-level multiprocessor and the like that combine two or more independent cores (called a die).

[3602] The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software on a server, client, firewall, gateway, hub, router, or other such computer and/or networking hardware. The software program may be associated with a server that may include a file server, print server, domain server, internet server, intranet server, cloud server, and other variants such as secondary server, host server, distributed server, and the like. The server may include one or more of memories, processors, computer readable transitory and/or non-transitory media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other servers, clients, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the server. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the server.

[3603] The processor may enable execution of multiple threads. These multiple threads may correspond to different programs. In various embodiments, a single program may be implemented as multiple threads by the programmer or may be decomposed into multiple threads by the processing hardware. The threads may be executed simultaneously to enhance the performance of the processor and to facilitate simultaneous operations of the application. A processor may be implemented as a packaged semiconductor die. The die includes one or more processing cores and may include additional functional blocks, such as cache. In various embodiments, the processor

may be implemented by multiple dies, which may be combined in a single package or packaged separately.

[3604] The networking hardware may include one or more interface circuits. In some examples, the interface circuit(s) may implement wired or wireless interfaces that connect, directly or indirectly, to one or more networks. Examples of networks include a cellular network, a local area network (LAN), a wireless personal area network (WPAN), a metropolitan area network (MAN), and/or a wide area network (WAN). The networks may include one or more of point-to-point and mesh technologies. Data transmitted or received by the networking components may traverse the same or different networks. Networks may be connected to each other over a WAN or point-to-point leased lines using technologies such as Multiprotocol Label Switching (MPLS) and virtual private networks (VPNs).

[3605] Examples of cellular networks include GSM, GPRS, 3G, 4G, 5G, LTE, and EVDO. The cellular network may be implemented using frequency division multiple access (FDMA) network or code division multiple access (CDMA) network. Examples of a LAN are Institute of Electrical and Electronics Engineers (IEEE) Standard 802.11-2020 (also known as the WIFI wireless networking standard) and IEEE Standard 802.3-2018 (also known as the ETHERNET wired networking standard). Examples of a WPAN include IEEE Standard 802.15.4, including the ZIGBEE standard from the ZigBee Alliance. Further examples of a WPAN include the BLUETOOTH™ wireless networking standard, including Core Specification versions 3.0, 4.0, 4.1, 4.2, 5.0, and 5.1 from the Bluetooth Special Interest Group (SIG). A WAN may also be referred to as a distributed communications system (DCS). One example of a WAN is the internet.

[3606] Storage hardware is or includes a computer-readable medium. The term computer-readable medium, as used in this disclosure, encompasses both nonvolatile storage and volatile storage, such as dynamic random access memory (DRAM). The term computer-readable medium only excludes transitory electrical or electromagnetic signals propagating through a medium (such as on a carrier wave). A computer-readable medium in this disclosure is therefore non-transitory, and may also be considered to be tangible.

[3607] Examples of storage implemented by the storage hardware include a database (such as a relational database or a NoSQL database), a data store, a data lake, a column store, and a data warehouse. Examples of storage hardware include nonvolatile memory devices, volatile memory devices, magnetic storage media, a storage area network (SAN), network-attached storage (NAS), optical storage media, printed media (such as bar codes and magnetic ink), and paper media (such as punch cards and paper tape). The storage hardware may include cache memory, which may be collocated with or integrated with processing hardware. Storage hardware may have read-only, write-once, or read/write properties. Storage hardware may be random access or sequential access. Storage hardware may be location-addressable, file-addressable, and/or content-addressable.

[3608] Examples of nonvolatile memory devices include flash memory (including NAND and NOR technologies), solid state drives (SSDs), an erasable programmable read-only memory device such as an electrically erasable programmable read-only memory (EEPROM) device, and a mask read-only memory device (ROM). Examples of volatile memory devices include processor

registers and random access memory (RAM), such as static RAM (SRAM), dynamic RAM (DRAM), synchronous DRAM (SDRAM), synchronous graphics RAM (SGRAM), and video RAM (VRAM). Examples of magnetic storage media include analog magnetic tape, digital magnetic tape, and rotating hard disk drive (HDDs). Examples of optical storage media include a CD (such as a CD-R, CD-RW, or CD-ROM), a DVD, a Blu-ray disc, and an Ultra HD Blu-ray disc.

[3609] Examples of storage implemented by the storage hardware include a distributed ledger, such as a permissioned or permissionless blockchain. Entities recording transactions, such as in a blockchain, may reach consensus using an algorithm such as proof-of-stake, proof-of-work, and proof-of-storage. Elements of the present disclosure may be represented by or encoded as non-fungible tokens (NFTs). Ownership rights related to the non-fungible tokens may be recorded in or referenced by a distributed ledger. Transactions initiated by or relevant to the present disclosure may use one or both of fiat currency and cryptocurrencies, examples of which include bitcoin and ether. Some or all features of hardware may be defined using a language for hardware description, such as IEEE Standard 1364-2005 (commonly called "Verilog") and IEEE Standard 1076-2008 (commonly called "VHDL"). The hardware description language may be used to manufacture and/or program hardware.

[3610] A special-purpose system may be distributed across multiple different software and hardware entities. Communication within a special-purpose system and between special-purpose systems may be performed using networking hardware. The distribution may vary across embodiments and may vary over time. For example, the distribution may vary based on demand, with additional hardware and/or software entities invoked to handle higher demand. In various embodiments, a load balancer may direct requests to one of multiple instantiations of the special purpose system. The hardware and/or software entities may be physically distinct and/or may share some hardware and/or software, such as in a virtualized environment. Multiple hardware entities may be referred to as a server rack, server farm, data center, etc.

[3611] Software includes instructions that are machine-readable and/or executable. Instructions may be logically grouped into programs, codes, methods, steps, actions, routines, functions, libraries, objects, classes, etc. Software may be stored by storage hardware or encoded in other hardware. Software encompasses (i) descriptive text to be parsed, such as HTML (hypertext markup language), XML (extensible markup language), and JSON (JavaScript Object Notation), (ii) assembly code, (iii) object code generated from source code by a compiler, (iv) source code for execution by an interpreter, (v) bytecode, (vi) source code for compilation and execution by a just-in-time compiler, etc. As examples only, source code may be written using syntax from languages including C, C++, JavaScript, Java, Python, R, etc. The computer executable code may be created using a structured programming language such as C, an object oriented programming language such as C++, or any other high-level or low-level programming language (including assembly languages, hardware description languages, and database programming languages and technologies) that may be stored, compiled or interpreted to run on one of the devices described in the disclosure, as well as heterogeneous combinations of processors, processor architectures, or

combinations of different hardware and software, or any other machine capable of executing program instructions. Computer software may employ virtualization, virtual machines, containers, dock facilities, portainers, and other capabilities. In example embodiments, methods described in the disclosure and combinations thereof may be embodied in computer executable code that, when executing on one or more computing devices, performs the steps thereof. In another aspect, the methods may be embodied in systems that perform the steps thereof and may be distributed across devices in a number of ways, or all of the functionality may be integrated into a dedicated, standalone device or other hardware. In another aspect, the means for performing the steps associated with the processes described in the disclosure may include any of the hardware and/or software described in the disclosure. All such permutations and combinations are intended to fall within the scope of the disclosure.

[3612] The elements depicted in the flow chart and block diagrams or any other logical component may be implemented on a machine capable of executing program instructions. Thus, while the foregoing drawings and descriptions set forth functional aspects of the disclosed systems, no particular arrangement of software for implementing these functional aspects should be inferred from these descriptions unless explicitly stated or otherwise clear from the context. Similarly, it will be appreciated that the various steps identified and described in the disclosure may be varied, and that the order of steps may be adapted to particular applications of the techniques disclosed herein. All such variations and modifications are intended to fall within the scope of this disclosure. As such, the depiction and/or description of an order for various steps should not be understood to require a particular order of execution for those steps, unless required by a particular application, or explicitly stated or otherwise clear from the context.

[3613] Software also includes data. However, data and instructions are not mutually-exclusive categories. In various embodiments, the instructions may be used as data in one or more operations. As another example, instructions may be derived from data. The functional blocks and flowchart elements in this disclosure serve as software specifications, which can be translated into software by the routine work of a skilled technician or programmer. Software may include and/or rely on firmware, processor microcode, an operating system (OS), a basic input/output system (BIOS), application programming interfaces (APIs), libraries such as dynamic-link libraries (DLLs), device drivers, hypervisors, user applications, background services, background applications, etc. Software includes native applications and web applications. For example, a web application may be served to a device through a browser using hypertext markup language 5th revision (HTML5).

[3614] Software may include artificial intelligence systems, which may include machine learning or other computational intelligence. For example, artificial intelligence may include one or more models used for one or more problem domains. When presented with many data features, identification of a subset of features that are relevant to a problem domain may improve prediction accuracy, reduce storage space, and increase processing speed. This identification may be referred to as feature engineering. Feature engineering may be performed by users or may only be guided by users. In various implementations, a machine learning system may computationally identify

relevant features, such as by performing singular value decomposition on the contributions of different features to outputs.

[3615] Examples of the models include recurrent neural networks (RNNs) such as long short-term memory (LSTM), deep learning models such as transformers, decision trees, support-vector machines, genetic algorithms, Bayesian networks, and regression analysis. Examples of systems based on a transformer model include bidirectional encoder representations from transformers (BERT) and generative pre-trained transformers (GPT). Training a machine-learning model may include supervised learning (for example, based on labelled input data), unsupervised learning, and reinforcement learning. In various embodiments, a machine-learning model may be pre-trained by their operator or by a third party. Problem domains include nearly any situation where structured data can be collected, and includes natural language processing (NLP), computer vision (CV), classification, image recognition, etc.

[3616] Some or all of the software may run in a virtual environment rather than directly on hardware. The virtual environment may include a hypervisor, emulator, sandbox, container engine, etc. The software may be built as a virtual machine, a container, etc. Virtualized resources may be controlled using, for example, a DOCKER container platform, a pivotal cloud foundry (PCF) platform, etc.

[3617] In a client-server model, some of the software executes on first hardware identified functionally as a server, while other of the software executes on second hardware identified functionally as a client. The identity of the client and server is not fixed: for some functionality, the first hardware may act as the server while for other functionality, the first hardware may act as the client. In different embodiments and in different scenarios, functionality may be shifted between the client and the server. In one dynamic example, some functionality normally performed by the second hardware is shifted to the first hardware when the second hardware has less capability. In various embodiments, the term "local" may be used in place of "client," and the term "remote" may be used in place of "server."

[3618] Some or all of the software may be logically partitioned into microservices. Each microservice offers a reduced subset of functionality. In various embodiments, each microservice may be scaled independently depending on load, either by devoting more resources to the microservice or by instantiating more instances of the microservice. In various embodiments, functionality offered by one or more microservices may be combined with each other and/or with other software not adhering to a microservices model.

[3619] Some or all of the software may be arranged logically into layers. In a layered architecture, a second layer may be logically placed between a first layer and a third layer. The first layer and the third layer would then generally interact with the second layer and not with each other. In various embodiments, this is not strictly enforced – that is, some direct communication may occur between the first and third layers.

FURTHER INFORMATION AND USE OF TERMS

[3620] The background description is presented simply for context, and is not necessarily well-understood, routine, or conventional. Further, the background description is not an admission of

what does or does not qualify as prior art. In fact, some or all of the background description may be work attributable to the named inventors that is otherwise unknown in the art.

[3621] While only a few embodiments of the disclosure have been shown and described, it will be obvious to those skilled in the art that many changes and modifications may be made thereunto without departing from the spirit and scope of the disclosure as described in the following claims. All patent applications and patents, both foreign and domestic, and all other publications referenced herein are incorporated herein in their entireties to the full extent permitted by law. While the disclosure has been disclosed in connection with the preferred embodiments shown and described in detail, various modifications and improvements thereon will become readily apparent to those skilled in the art. Accordingly, the spirit and scope of the disclosure is not to be limited by the foregoing examples, but is to be understood in the broadest sense allowable by law.

[3622] The detailed description includes specific examples for illustration only, and not to limit the disclosure or its applicability. The examples are not intended to be an exhaustive list, but instead simply demonstrate possession by the inventors of the full scope of the currently presented and envisioned future claims. Variations, combinations, and equivalents of the examples are within the scope of the disclosure. No language in the specification should be construed as indicating that any non-claimed element is essential or critical to the practice of the disclosure.

[3623] While the foregoing written description enables one skilled to make and use what is considered presently to be the best mode thereof, those skilled in the art will understand and appreciate the existence of variations, combinations, and equivalents of the specific embodiment, method, and examples herein. The disclosure should therefore not be limited by the above-described embodiment, method, and examples, but by all embodiments and methods within the scope and spirit of the disclosure.

[3624] Each publication referenced in this disclosure, including foreign and domestic patent applications and patents, is hereby incorporated by reference in its entirety as if fully set forth herein.

[3625] The terms "comprising," "with," "including," and "containing" are to be construed as open-ended terms (i.e., meaning "including, but not limited to") unless otherwise noted. Unless otherwise specified, the terms "comprising," "having," "with," "including," and "containing," and their variants, are open-ended terms, meaning "including, but not limited to."

[3626] The term "exemplary" simply means "example" and does not indicate a best or preferred example. The use of any and all examples, or exemplary language (e.g., "such as") provided herein, is intended merely to better illuminate the disclosure and does not pose a limitation on the scope of the disclosure unless otherwise claimed.

[3627] The use of the terms "a," "an," "the," and similar referents in the context of describing the disclosure (especially in the context of the following claims) is to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context.

[3628] Recitations of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually

recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context.

[3629] The phrase "at least one of A, B, and C" should be construed to mean a logical (A OR B OR C), using a non-exclusive logical OR, and should not be construed to mean "at least one of A, at least one of B, and at least one of C."

[3630] The term "set" may include a set with a single member. The term "set" does not necessarily exclude the empty set — in other words, in some circumstances a "set" may have zero elements. The term "non-empty set" may be used to indicate exclusion of the empty set — that is, a non-empty set must have one or more elements. The term "subset" does not necessarily require a proper subset. In other words, a "subset" of a first set may be coextensive with (equal to) the first set. Further, the term "subset" does not necessarily exclude the empty set — in some circumstances a "subset" may have zero elements.

[3631] Physical (such as spatial and/or electrical) and functional relationships between elements (for example, between modules, circuit elements, semiconductor layers, etc.) are described using various terms. Unless explicitly described as being "direct," when a relationship between first and second elements is described, that relationship encompasses both (i) a direct relationship where no other intervening elements are present between the first and second elements and (ii) an indirect relationship where one or more intervening elements are present between the first and second elements. Example relationship terms include "adjoining," "transmitting," "receiving," "connected," "engaged," "coupled," "adjacent," "next to," "on top of," "above," "below," "abutting," and "disposed."

[3632] Although each of the embodiments is described above as having certain features, any one or more of those features described with respect to any embodiment of the disclosure can be implemented in and/or combined with features of any of the other embodiments, even if that combination is not explicitly described. In other words, the described embodiments are not mutually exclusive, and permutations of multiple embodiments remain within the scope of this disclosure.

[3633] One or more elements (for example, steps within a method, instructions, actions, or operations) may be executed in a different order (and/or concurrently) without altering the principles of the present disclosure. Unless technically infeasible, elements described as being in series may be implemented partially or fully in parallel. Similarly, unless technically infeasible, elements described as being in parallel may be implemented partially or fully in series.

[3634] While the disclosure describes structures corresponding to claimed elements, those elements do not necessarily invoke a means plus function interpretation unless they explicitly use the signifier "means for." Unless otherwise indicated, recitations of ranges of values are merely intended to serve as a shorthand way of referring individually to each separate value falling within the range, and each separate value is hereby incorporated into the specification as if it were individually recited.

[3635] While the drawings divide elements of the disclosure into different functional blocks or action blocks, these divisions are for illustration only. According to the principles of the present

disclosure, functionality can be combined in other ways such that some or all functionality from multiple separately-depicted blocks can be implemented in a single functional block; similarly, functionality depicted in a single block may be separated into multiple blocks. Unless explicitly stated as mutually exclusive, features depicted in different drawings can be combined consistent with the principles of the present disclosure.

[3636] In the drawings, reference numbers may be reused to identify identical elements or may simply identify elements that implement similar functionality. Numbering or other labeling of instructions or method steps is done for convenient reference, not to indicate a fixed order. In the drawings, the direction of an arrow, as indicated by the arrowhead, generally demonstrates the flow of information (such as data or instructions) that is of interest to the illustration. For example, when element A and element B exchange a variety of information but information transmitted from element A to element B is relevant to the illustration, the arrow may point from element A to element B. This unidirectional arrow does not imply that no other information is transmitted from element B to element A. As one example, for information sent from element A to element B, element B may send requests and/or acknowledgements to element A.

CLAIMS

What is claimed is:

1. An industrial technology stack for an industrial environment, the industrial technology stack comprising:

a set of computational resources; and

a set of layers executed by the set of computational resources, the set of layers including,

a governance layer,

an enterprise layer,

an offering layer,

a transaction layer,

an operations layer,

a network layer,

a data layer, and

a resource layer.

2. The industrial technology stack of claim 1, wherein the industrial technology stack performs at least one of,

enabling devices of the industrial environment to cooperatively and interoperatively discover and assign resources,

orchestrating an application of computational resources to industrial resources and uses,

monitoring and manage industrial processes, or

maintaining at least one resource of the industrial environment.

3. The industrial technology stack of claim 1, wherein the governance layer monitors resources of the industrial environment to ensure compliance with governance policies of the industrial environment, and the governance policies include at least one of,

conserving consumption of resources of the industrial environment,

improving efficiency of industrial processes of the industrial environment,

promoting productivity of employees of the industrial environment,

reducing greenhouse gas emissions by the industrial environment, or

ensuring compliance with at least one law, regulation, moral principle, ethical obligation, or organizational objective.

4. The industrial technology stack of claim 1, wherein the governance layer includes at least one of,

artificial intelligence models for embedded policies and governance,

policy automation,

safety and compliance automation,

reporting automation, or

automated governance of the industrial environment.

5. The industrial technology stack of claim 1, wherein the enterprise layer performs at least one of,

allocating resources of an industrial enterprise to industrial processes and/or equipment of the industrial environment, or

selecting, analyzing, and/or adapting an organizational structure of the industrial environment to entities within the industrial environment.

6. The industrial technology stack of claim 1, wherein the enterprise layer includes at least one,

artificial intelligence models for contextual simulation and forecasting,

executive digital twins,

IIoT platforms,

big data and analytics, or

advanced operational control over the industrial environment.

7. The industrial technology stack of claim 1, wherein the offering layer performs at least one of,

selecting and/or describing external contracts or opportunities for participation by the industrial environment,

formulating, analyzing, and/or releasing external offerings for participation of other enterprises, or

selecting, describing, formulating, analyzing, and/or releasing internal offerings among at least two entities of the industrial environment.

8. The industrial technology stack of claim 1, wherein the offering layer includes at least one of,

expert systems,

generative AI,

industrial metaverse simulation of the industrial environment,

systems integration of the industrial environment,

industrial control systems of the industrial environment, or

industrial automation of the industrial environment.

9. The industrial technology stack of claim 1, wherein the transaction layer formulates, negotiates, executes, fulfills, and/or reports transactions related to one or more offerings of the offering layer, and the transactions include at least one of,

a purchase of one or more assets for the industrial environment,

a sale of one or more assets of the industrial environment,

a mortgage, lease, and/or loan of one or more resources of the industrial environment, or

a conveyance, transport, and/or delivery of a physical or digital asset of the industrial environment.

10. The industrial technology stack of claim 1, wherein the transaction layer includes at least one of,

artificial intelligence models for discovery, generation, and optimization of transactions,

smart contracts and distributed ledgers,

a transaction-related system-of-systems,

automated transaction orchestration, or

an automated industrial marketplace.

11. The industrial technology stack of claim 1, wherein the operations layer performs at least one of,

executing steps of an industrial manufacturing, refining, materials processing, inspection, assembly, and/or delivery process,

monitoring a set of resources and/or operations of the industrial environment,

assessing a set of resources and/or operations of the industrial environment,

diagnosing a set of resources and/or operations of the industrial environment,

maintaining a set of resources and/or operations of the industrial environment,

repairing a set of resources and/or operations of the industrial environment,

cleaning a set of resources and/or operations of the industrial environment,

generating a set of resources of the industrial environment,

organizing a set of resources of the industrial environment,

interconnecting a set of resources of the industrial environment, or

maintaining a set of resources of the industrial environment.

12. The industrial technology stack of claim 1, wherein the operations layer includes at least one of,

artificial intelligence models for routing, control, optimization, and/or generation of resources of the industrial environment,

automated monitoring and maintenance platforms,

asset and infrastructure management,

robotics and process automation, or

converged, AI-based IT/OT workflow orchestration.

13. The industrial technology stack of claim 1, wherein the network layer provides connectivity and/or communicating among a set of components of the industrial environment, wherein the set of components of the industrial environment includes at least one of,

an industrial machine,

a computer,

a server,

a workstation,

an industrial sensor,

a personal communication device, or

a software component of the industrial technology stack.

14. The industrial technology stack of claim 1, wherein the network layer includes at least one of,

artificial intelligence models for adaptive networking,

edge and/or cloud communication,

cellular, Wi-Fi, ORAN, and/or Bluetooth communication protocols,

Industrial Internet-of-Things (IIoT)-related communication protocols, or

intelligence edge, advanced networking, and/or IIoT-in-a-Box architectures.

15. The industrial technology stack of claim 1, wherein the data layer performs at least one of,

providing data storage capabilities of the industrial environment,
managing data storage capabilities of the industrial environment,
storing data of the industrial environment,
curating data of the industrial environment,
indexing data of the industrial environment,
searching data of the industrial environment,
querying data of the industrial environment,
structuring data of the industrial environment,
providing data of the industrial environment,
refreshing data of the industrial environment,
compiling data of the industrial environment,
transmitting data of the industrial environment,
altering data of the industrial environment,
monitoring data of the industrial environment, or
deleting data of the industrial environment.

16. The industrial technology stack of claim 1, wherein the data layer includes at least one of,

artificial intelligence models for sensor fusion,
sensor and/or wearable devices of the industrial environment,
mobile data collectors of the industrial environment,
application programming interfaces of the industrial environment, or
context-aware industrial IoT and sensor fusion.

17. The industrial technology stack of claim 1, wherein the resource layer performs at least one of,

acquiring a resource of the industrial environment,
managing a resource of the industrial environment,
monitoring a resource of the industrial environment,
controlling a resource of the industrial environment,
transporting a resource of the industrial environment, or
maintaining a resource of the industrial environment.

18. The industrial technology stack of claim 1, wherein the resource layer includes at least one of,

artificial intelligence models for resource optimization,
advanced computation,
switches and/or multiplexers,
vibration or other sensor data collection, or
machine-level data collection and resource optimization.

19. A system for visual search in a digital twin environment, comprising:

a visual search engine embedded within a digital twin;
an interface configured to receive an image-based query; and
one or more processors configured to:
process the image-based query and data structures supporting the digital twin to identify matching elements in the digital twin,
present the matching elements in a visual representation of the digital twin, and
provide additional data associated with the matching elements.

20. The system of claim 19, wherein the image-based query comprises at least one of:
a pasted image,
image metadata,
a user-created drawing, or
a set of keywords describing an image.

21. The system of claim 19, wherein presenting the matching elements comprises at least one of:
listing the matching elements in a menu bar, or
highlighting the matching elements in the visual representation through changes in font, color, or brightness.

22. The system of claim 19, wherein the processors are further configured to utilize artificial intelligence to perform at least one of,
parsing queries,
recognizing visual elements,
classifying visual elements,
matching images with image content,
correlating images with image content,
ranking results based on likelihood of relevance, or
optimizing a presentation of results.

23. A system for automatically constructing a digital twin environment, comprising one or more processors configured to:
receive an arbitrary set of photographs taken over time,
utilize artificial intelligence systems to link common elements within the photographs,
manipulate image elements to provide a stitched view of the digital twin environment, and
add information to the digital twin from other data sources.

24. The system of claim 23, further comprising a blockchain configured to,
capture image information and other information about the environment,
support a distributed ledger of transactions related to the environment,
provide interfaces for controlling access to the information, and
enable private and public portions of the information.

25. A system for managing decision-making in a digital twin environment, comprising:
a digital twin platform configured to display states and activities of entities; and

one or more processors configured to:

simulate, based on data for the entities, states and activities processed by the digital twin, a set of outcomes predicted to result upon deployment of a set of human and artificial intelligence systems in decision making workflows;

provide a planning system for deployment of artificial intelligence systems and required resources; and

predict outcomes of deployment of multiple distinctive sets of artificial intelligence systems and resource configurations.

26. The system of claim 25, wherein the processors are further configured to display training metrics for at least one of,

humans available to perform decision making tasks,

hybrid human-AI systems available to perform decision making tasks, or

AI systems available to perform decision making tasks.

27. The system of claim 25, wherein the processors are further configured to display performance metrics for at least one of,

humans available to perform decision making tasks,

hybrid human-AI systems available to perform decision making tasks, or

AI systems available to perform decision making tasks.

28. The system of claim 25, wherein the processors are further configured to display comparative performance metrics among at least one of,

individual humans,

human-led enterprises,

crowds of humans,

human-AI combination systems, or

AI systems.

29. The system of claim 25, further comprising an intelligent agent configured to: automatically discover available systems among human systems, combined human-AI systems and standalone artificial intelligence systems capable of performing desired functions, and recommend deployment configuration parameters based on situational factors.

30. The system of claim 25, wherein the situational factors comprise at least one of,

availability factors for human and artificial intelligence resources,

capability factors for human and artificial intelligence resources,

resource availability factors for enabling human and artificial intelligence resources, or

contextual and environmental factors.

31. The system of claim 25, further comprising an interface system configured to,

designate users as trainers for creation of artificial intelligence systems represented in the digital twin, and

designate users as supervisors for artificial intelligence systems represented in the digital twin.

32. The system of claim 25, wherein the processors are further configured to distribute decision making authority based on a framework including at least one of,

a hierarchical framework,

a rules-based framework,

a simulation framework,

an enterprise planning framework,

an algorithmic framework,

a principles-based framework,

a collaborative framework,

a peer-to-peer framework, or

a competitive framework.

33. The system of claim 25, further comprising a smart contract system embedded in the digital twin configured to,

set terms and conditions for liability resulting from system outputs, and

process input data to produce outputs reflecting transactional terms.

34. The system of claim 25, wherein the processors are further configured to enable at least one of,

enabling incremental change,

enabling experimentation,

enabling reversible implementation,

enabling graceful migration among humans, human-AI mixes and autonomous AI systems, or

enabling rapid adjustment based on performance capabilities, contextual factors and environmental factors.

35. A system for managing data in a digital twin environment, comprising:

a pipeline of IoT and edge devices;

high QoS networks configured for adaptive networking; and

one or more processors configured to,

collect granular, real-time sensor data about entities in an operating environment,

process the sensor data for big data analytics, and

process the sensor data for machine learning.

36. The system of claim 35, further comprising data architectures configured to perform at least one of,

adapting visual representations for human cognitive processing based on user parameters, or

preparing and stage data for artificial intelligence systems through at least one of,

data cleansing,

deduplication,

generation of synthetic data,

entity resolution,

normalization, or
generation of embeddings.

37. The system of claim 35, further comprising computational architectures configured to adapt processing using a mix of:
cloud processing systems;
edge processing systems;
on-device processing systems; and
quantum computing systems.

38. A system for managing artificial intelligence deployment, comprising:
one or more processors configured to:
evaluate trust factors including:
stakes involved;
training data bias;
governance compliance;
performance metrics;
data quality and availability;
connectivity quality and availability;
computational resource availability; and
environmental parameters; and
determine extent of permitted AI system autonomy based on the trust factors.

39. The system of claim 38, wherein the processors are configured to perform at least one of,
permitting AI system control during normal operations; or
modifying AI system control permissions based on at least one of,
environmental catastrophes,
market swings, or
other significant environmental shifts.

40. A system for training artificial intelligence in a digital twin environment, comprising:
interfaces configured to receive from human experts at least one of,
tagged data sets for AI classification,
code, text, images, audio and video for training generative AI,
preference indicators for recommendation systems,
system configurations for configuration recommendations, or
decision examples for similar contexts, and
processors configured to train AI systems using the received data.

41. The system of claim 40, wherein the processors are configured to,
designate first users to provide input data for initial AI system training,
designate second users to supervise AI system outputs, and
enable overlap between the first and second users.

42. A system for managing enterprise workflows, comprising:
systems integration configured to integrate,

individual activities,
group activities,
enterprise activities, and
business ecosystem activities;

data and sensor fusion configured to track outcomes using at least one of,
environmental data,
market data, or
transaction data; and

AI-driven stakeholder interfaces configured to perform at least one of,
filtering information,
prioritizing information, or
distributing information for planning, simulation, decision making and operational action.

43. The system of claim 42, wherein the processors are configured to perform at least one of,
enabling collaboration,
distributing decision making to combinations of at least one of,
human beings,
artificial intelligence systems, or
other systems, or
removing fear of closing automation loops through at least one of,
incremental change capabilities,
experimentation capabilities, or
easily reversible implementation.

44. The system of claim 42, further comprising a simulation environment configured to,
use historical data,
use current data,
predict outcomes from deployment configurations, and
provide a planning environment for resource deployment scenarios.

45. A computer-implemented system for managing operations of artificial intelligence systems, comprising:
a processor; and
a memory storing instructions that, when executed by the processor, cause the system to:
store a set of operations, wherein each operation of the set of operations is associated with at least one artificial intelligence system of a set of artificial intelligence systems; and
initiate at least one operation of the stored set of operations with at least one artificial intelligence system of the set of artificial intelligence systems that is associated with the at least one operation.

46. The system of claim 45, wherein the instructions further cause the system to:

generate at least one artificial intelligence system of the set of artificial intelligence systems;

store the at least one generated artificial intelligence system; and

deploy the at least one generated artificial intelligence system to at least one device.

47. The system of claim 45, wherein the instructions further cause the system to:

generate a training data set for training at least one artificial intelligence system;

train the at least one artificial intelligence system using the generated training data set; and

verify a performance of the trained artificial intelligence system.

48. The system of claim 45, wherein the instructions further cause the system to:

monitor a performance of at least one artificial intelligence system during execution of at least one operation;

analyze the monitored performance; and

update the at least one artificial intelligence system based on the analyzed performance.

49. The system of claim 45, wherein the instructions further cause the system to:

orchestrate operation of at least a portion of the set of artificial intelligence systems to fulfill a set of requests;

schedule invocation of artificial intelligence systems based on processing loads and throughput; and

perform load-balancing of requests to invoke artificial intelligence systems.

50. The system of claim 45, wherein the instructions further cause the system to:

adapt at least one artificial intelligence system for deployment to a particular device;

package resources of the adapted artificial intelligence system for deployment; and

transmit the packaged resources to the particular device.

51. The system of claim 45, wherein the instructions further cause the system to:

aggregate two or more artificial intelligence systems into a hybrid system that combines heterogeneous processing capabilities;

connect at least one output of a first artificial intelligence system to at least one input of a second artificial intelligence system; and

condition processing of data by the first artificial intelligence system based on processing of the data by the second artificial intelligence system.

52. The system of claim 45, wherein the instructions further cause the system to:

govern operation of at least one artificial intelligence system based on legal, social, practical, or regulatory policies;

monitor invocations of the at least one artificial intelligence system to detect unauthorized uses; and

monitor input to and output from the at least one artificial intelligence system to detect unauthorized content.

53. The system of claim 45, wherein the instructions further cause the system to:

log activities of the artificial intelligence systems during performance of operations;

correlate activities of various artificial intelligence systems to determine causal relationships; and

store the logged activities in one or more event logs.

54. The system of claim 45, wherein the set of operations includes:

operations for generating artificial intelligence models;

operations for evaluating artificial intelligence models;

operations for deploying artificial intelligence models; and

operations for invoking artificial intelligence models to perform various operations.

55. A computer-implemented system for managing operations of artificial intelligence systems, comprising:

a processor; and

a memory storing instructions that, when executed by the processor, cause the system to:

generate synthetic data to augment training data sets;

evaluate the synthetic data for hallmarks of synthetic data; and

adapt augmentation techniques to reduce artifacts of synthetic data.

56. The system of claim 55, wherein the instructions further cause the system to:

detect reduced performance of an artificial intelligence system;

initiate retraining of the artificial intelligence system using the augmented training data sets; and

verify that a retrained artificial intelligence system addresses the reduced performance.

57. The system of claim 55, wherein the instructions further cause the system to:

generate a plurality of generative artificial intelligence systems that produce content to be combined;

generate primary content using a first artificial intelligence system; and

generate supplemental content for the primary content using a second artificial intelligence system.

58. The system of claim 55, wherein the instructions further cause the system to:

invoke a first artificial intelligence system to generate content;

invoke a second artificial intelligence system to review the generated content; and

alter the generated content to address issues indicated in the review.

59. The system of claim 55, wherein the instructions further cause the system to:

monitor social, political, economic, scientific, and technical events;

extrapolate effects of current events on operation of artificial intelligence systems; and

execute actions to promote value, security, and longevity of artificial intelligence systems.

60. The system of claim 55, wherein the instructions further cause the system to:

monitor invocations of artificial intelligence systems to detect malicious use;

detect attempts to circumvent governance mechanisms; and

update policies and governance mechanisms to prevent detected circumvention attempts.

61. The system of claim 55, wherein the instructions further cause the system to:

generate an artificial intelligence agent configured to monitor conditions;

configure the artificial intelligence agent to learn policies based on experience; and
configure the artificial intelligence agent to reflect on policies to identify synergies and conflicts.

62. The system of claim 55, wherein the instructions further cause the system to:
determine hyperparameters of artificial neural networks;
generate artificial neural networks according to determined hyperparameters; and
store resources of generated artificial neural networks in a data store.

63. The system of claim 55, wherein the instructions further cause the system to:
partition training data sets into training, validation and test data sets;
analyze training data sets for training issues; and
request collection of additional data to address identified training issues.

64. The system of claim 55, wherein the instructions further cause the system to:
apply supervision to inputs of generative artificial intelligence systems;
apply supervision to processing of generative artificial intelligence systems; and
apply supervision to outputs of generative artificial intelligence systems based on rules, guidelines, heuristics or principles.