REAL-TIME SAVING OF HIGH-RESOLUTION IMAGES ON COMMERCIAL AUTOMATED VISUAL INSPECTION (AVI) SYSTEMS

FIELD OF THE DISCLOSURE

[0001] The present disclosure generally relates to systems and methods for automated visual inspection (AVI) in a production line, and, more particularly, to saving high-resolution images in real-time.

BACKGROUND

[0002] As the resolution and frame rates of industrial cameras have increased substantially, systems for saving image data in real time have struggled to keep pace. For example, the imaging system may produce high-resolution images exceeding 1MB per image. In addition to the higher image resolution of contemporary AVI systems, each product can produce dozens of different images at a single camera station. Accordingly, the collection of image data that must be processed within the framerate may exceed 150MB of image data per container. To allow a high-speed AVI machine to operate at full capacity, image retention capabilities are generally disabled for production runs.

[0003] One technical problem with saving images arises from the narrow temporal window available for inspection on a high-speed production line. Typically, a 300 parts-per-minute system might have 200msec to perform inspection. Even with fast industrial computers having optimized random access memory (RAM) and powerful graphics processing units (GPUs), present technology may not be sufficient to perform both inspection of image data and saving of the images in substantially real-time. For example, the conventional processing techniques of the collection of image data may utilize 190msec of the 200msec window, leaving very little time for saving the image data. For this reason, in general, conventional AVI systems are not able to provide comprehensive, robust real-time image saving capabilities or must reduce the resolution of the saved image data to lower the file size for the collection of image data to a size that can still be saved within the remaining time of the 200msec window. This size reduction results in lower accuracy when visual inspection analyses are applied to the saved image data.

[0004] The systems and methods disclosed herein provide solutions to these problems and others.

SUMMARY

[0005] The following relates to AVI systems and methods for automated inspection of products and saving image data and inspection metadata, such as image acquisition metadata, image processing metadata, and system metadata, in a production environment. As described herein, the disclosed systems and methods relate to a hardware configuration that enables the inspection and saving of image data in real-time. This can be achieved without or at least without a substantial reduction of the resolution of the original image data required to be able to save the image data within the processing window, thereby resulting is more accurate visual inspection analyses in real-time. Depending on the embodiment, reference herein to a "substantial" reduction in the resolution of the original image data may be a reduction of the resolution of one or more of the component images of a container by

75%, 50%, 35%, etc. and/or the exclusion of captured images (e.g., 75%, 50%, 35%, etc. of the images) of the container captured by the various imaging units of the disclosed AVI systems. In the context of this disclosure, real-time refers to the ability to both analyze and save the image data within a window of processing time for individual products. For example, if the AVI system has a throughput of 300 products per minute, the window of processing time may be about 200msec per product (depending on the particular components utilized in the AVI system).

[0006] Conventional AVI systems typically utilize the same hardware for both performing the AVI techniques and saving the image data. As AVI techniques have become more sophisticated and digital cameras produce higher resolution images (and thus larger sets of image data), the AVI analysis utilizes a significant amount of the processing resources within the processing window and may not provide sufficient time to save the image data by the time the image analysis is complete. In one scenario, the processing for the AVI analyses utilizes 95% of the processing window, leaving too little time to save the image data.

[0007] The disclosed AVI systems and methods may save the image data at a full production cycle speed without impacting core product inspection functionality by parallelizing the AVI analysis and the saving of the image data across different hardware systems. The image data may be saved without or at least without a substantial reduction of the resolution of the original image data allowing real-time visual inspection as well as, for example, preparing training data sets for machine learning. As a result, the AVI system is able to perform the AVI analyses without utilizing processing resources needed to save the underlying image data. Additionally, if additional hardware units are included in the AVI system, the parallelization techniques enable additional AVI techniques to be applied within the processing window. As a result, the disclosed embodiments enable the performance of additional AVI techniques within the processing window, thereby enabling more thorough inspection of the individual products.

[0008] In one aspect, an AVI system may be provided. The AVI system may include: (1) an imaging unit for generating image data of a product; (2) a hardware image splitter comprising an input port, a first output port, and a second output port, wherein the imaging unit is coupled to the input port; (3) a first image analysis hardware unit, coupled to the first output port, comprising one or more processors and a memory storing computer-readable instructions, wherein the instructions cause the first image analysis hardware unit to: (i) receive the image data from the hardware image splitter, (ii) analyze the image data for product defects using a first image inspection analysis, and (iii) generate first image processing metadata based on the first image inspection analysis; (4) an image storage hardware unit, coupled to the second output port, comprising one or more processors, a buffer, and a memory storing computer-readable instructions, wherein the instructions cause the image saving computer to: (i) store the image data from the hardware image splitter into the buffer, (ii) receive the first image processing metadata, and (iii) associate inspection metadata with the image data, wherein the inspection metadata includes the first image processing metadata.

[0009] In one aspect, an AVI-based method may be provided. The method may include: (1) generating, using an imaging unit, image data of a product; (2) receiving, at a first image analysis hardware unit, the image data from a hardware image splitter; (3) analyzing, using the first image analysis hardware unit, the image data for product defects using a first image inspection analysis; (4) generating, using the first image analysis hardware unit, the first image processing metadata based on the analysis; (5) receiving, at an image storage hardware unit, the image data from the hardware image splitter; (6) storing, using the image storage hardware unit, the image data into a buffer; (7) receiving, at the image storage hardware unit, the first image processing metadata; and (8) associating, using the image storage hardware unit, inspection metadata with the image data, wherein the inspection metadata includes the first image processing metadata.

[0010] Advantages will become more apparent to those of ordinary skill in the art from the following description of the preferred embodiments, which are shown and described by way of illustration. As will be realized, the present embodiments may be capable of other and different embodiments, and their details are capable of modification in various respects. Accordingly, the drawings and description are to be regarded as illustrative in nature and not as restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The figures described below depict various aspects of the system and methods disclosed therein. It should be understood that each figure depicts an embodiment of a particular aspect of the disclosed system and methods, and that each of the figures is intended to accord with a possible embodiment thereof. Further, wherever possible, the following description refers to the reference numerals included in the following figures, in which features depicted in multiple figures are designated with consistent reference numerals.

[0012] There are shown in the drawings arrangements which are presently discussed, it being understood, however, that the present embodiments are not limited to the precise arrangements and instrumentalities shown, wherein:

[0013] Figure 1 illustrates an exemplary AVI environment in which the techniques disclosed herein may be implemented, in accordance with various embodiments disclosed herein.

[0014] Figure 2 illustrates an exemplary image information and inspection metadata database table, in accordance with various embodiments disclosed herein.

[0015] Figure 3 illustrates an exemplary graphics file comprising inspection metadata, in accordance with various embodiments disclosed herein.

[0016] Figure 4 illustrates an exemplary timing diagram, in accordance with various embodiments disclosed herein.

[0017] Figure 5 illustrates an exemplary AVI method for image inspection and storage, in accordance with various embodiments disclosed herein.

[0018] The figures depict preferred embodiments for purposes of illustration only. Alternative embodiments of the systems and methods illustrated herein may be employed without departing from the principles of the invention described herein.

DETAILED DESCRIPTION

[0019] The various concepts introduced above and discussed in greater detail below may be implemented in any of numerous ways, and the described concepts are not limited to any particular manner of implementation. Examples of implementations are provided for illustrative purposes. Although examples described herein generally relate to pharmaceutical manufacturing processes and the visual inspection of pharmaceutical products, the techniques disclosed herein may be applied to other visual inspection processes.

[0020] Figure 1 depicts an exemplary automated visual inspection (AVI) environment 100 in which the techniques disclosed herein may be implemented, according to some aspects. The AVI environment 100 may include an imaging unit 110, products 120, a hardware image splitter 130, one or more image analysis hardware units 140A-140N, an image storage hardware unit 150, cloud storage 160, and system controller unit 170. One or more components of the AVI environment 100 may be located in a manufacturing facility, laboratory, fabrication plant, or other suitable environment for inspection of products.

[0021] The products 120 may comprise a plurality of manufactured, assembled, or fabricated items, such as circuit boards, semiconductor wafers, containers, etc. In pharmaceutical contexts, the products 120 may be a vial, a cartridge, a syringe, a vessel, etc. The products 120 may be transported through a production line by one or more types of automated conveyance equipment, such as conveyor belts, star tables, rotators, manipulators, arms, and so on. The production line in the AVI environment 100 may operate at a rate of 200 products per minute, 300 products per minute, 400 products per minute, 500 products per minute, 600 products per minute, 800 products per minute, 1000 products per minute, 2000 products per minute, and so on.

[0022] The system controller unit 170 may include one or more processors, memories, network interfaces, bus interfaces, and buffers. The one or more processors may be a programmable logic controller (PLC) that executes software instructions stored in a memory unit to perform some or all of the functions of the image storage hardware unit 150 as described herein. Alternatively, some of the processors may be other types of processors (e.g., application-specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), graphics processing units (GPUs), etc.), and some of the functionality of the image storage hardware unit 150 as described herein may instead be implemented, in part or in whole, by such hardware. The memory units may include one or more physical memory devices or units containing volatile and/or non-volatile memory.

[0023] The system controller unit 170 may measure and/or control operation of the AVI environment 100, including, for example, operation of one or more conveyors configured to move the products 120 (e.g., at a specified rate and/or in a specified direction of motion), one or more imaging units 110, and/or other equipment of the AVI environment 100 (including equipment not depicted by Figure 1) in accordance with an operating procedure. For example, the procedure may include a set of parameters (and, in some embodiments, timings associated therewith) that are interpretable by the system controller unit 170 to cause the equipment of the AVI environment 100 implement the indicated actions. For example, the procedure may indicate to system controller unit 170 how to control a light intensity associated with a light source proximate to an imaging unit 110, a frame rate associated with an imaging unit 110, a speed and/or pattern of a conveyor elements, and so on. In some embodiments, the system controller unit 170 may generate system metadata 172, which may comprise the one or more of the control settings indicated by the procedure. In AVI systems that do not rely on procedure files, the system metadata 172 may include comparable control parameters relied upon by a corresponding system controller unit.

[0024] The imaging unit 110 may comprise one or more digital cameras and/or image sensors having fields of view (FOVs) oriented at an imaging area associated with the production line. The imaging unit 110 may be a charge-coupled device (CCD) or a complementary metal-oxide semiconductor (CMOS) digital camera. Some example imaging units 110 that may be used in an AVI system include the Baumer VLXT-240M.I and LXC-250M, Vieworks VC-127MX2 or VP-103MC, or SVS-Vistek HR120. Accordingly, the individual images captured by the cameras and/or image sensors of the imaging unit 110 may have a resolution of at least 1.5 megapixels, at least 2 megapixels, at least 5 megapixels, at least 10 megapixels, at least 15 megapixels, at least 21 megapixels, or at least 26 megapixels, and so on.

[0025] The imaging unit 110 may generate image data 112 depicting the products 120 as the products 120 are carried through a production line. In the illustrated scenario, the imaging unit 110 is configured to capture image data 112 as the products 120 are carried by a conveyor. The cameras and/or image sensors of the imaging unit 110 may be positioned to have a field of view (FOV) oriented towards the products 120. In some embodiments, the cameras and/or image sensors of the imaging unit 110 may be positioned at different angles with respect to the products 120 in order to capture image data 112 from a plurality of perspectives. Similarly, in some embodiments, the production line is configured to rotate the products 120 within the FOV of the cameras and/or image sensors of the imaging unit 110 to capture image data 112 of the products 120 from a plurality of different rotational angles (e.g., ever 15º, every 30º, every 60º, and so on). In embodiments where the imaging unit 110 captures multiple images of each product 120, the total number of megapixels for the image data 112 associated with each product 120 may be at least 100 megapixels, at least 200 megapixels, at least 400 megapixels, at least 800 megapixels, or at least 1.2 gigapixels, and so on. Accordingly, even if one skilled in the art were to crop the images to reduce the overall amount of megapixels for the image data 112, there still remains a significant volume of image data such that conventional systems may

struggle to process and save the image data 112 in the absence of implementing the disclosed techniques. Accordingly, the AVI systems and methods described herein allow the image data to be saved without or at least without a substantial reduction of the resolution of the original image data.

[0026] The imaging unit 110 may be configured to compile the various images captured by cameras and/or image sensors of a given product 120 into a single set of image data 112 for the product 120. For example, the imaging unit 110 may compile each image generated by the one or more cameras and/or image sensors within a predefined time period in which a product 120 passes through the imaging region of the production line. In some embodiments, the imaging unit 110 compiles 25 images, 100 images, 150 images, or 200 images, and so on into the image data 112. In some embodiments, the compiled set of image data 112 is over 100MB in size, over 150MB in size, over 250MB in size, or over 400MB in size, and so on. As described above, conventional AVI systems may be unable to process image data of such sizes in real time for an AVI system that has a throughput of over 300 products per minute. Accordingly, the AVI systems and methods described herein allow the image data to be saved without or at least without a substantial reduction of the resolution of the original image data.

[0027] After capturing and/or compiling the image data 112, the imaging unit 110 may then label the image data 112 with any unique identifier that indicates the particular product 120 depicted in the image data 112. For example, the imaging unit 110 may save the image data 112 with a file name that indicates a lot number, a batch number, a product identifier, and/or any other type of identifier that is useful in identifying particular products 120. In some embodiments, the imaging unit 110 may include a memory that stores identifiers and/or a sequence of identifiers to be able to apply the appropriate identifier to the set of image data 112. As another example, the imaging unit 110 may save the image data 112 with a file name that indicates a time at which the image data 112 was captured. In these examples, another component (e.g., an image analysis hardware unit 140A) may identify a product identifier on an outer surface of the product 120 and correlate the time information with the detected identity of the product 120.

[0028] In some embodiments, the imaging unit 110 may generate image acquisition metadata 114. The image acquisition metadata 114 may include information about the image data 112, such as a timestamp when the image data 112 was captured, a gain applied by the imaging unit 110, a shutter speed/acquisition time setting of the imaging unit 110, a resolution of the image data 112, etc.

[0029] A hardware image splitter 130 may comprise a combination of hardware or firmware for splitting the input image data 112 into a plurality of image data 112 outputs. The hardware image splitter 130 may comprise an input port and a plurality (e.g., four) of output ports. The hardware image splitter 130 may receive the image data 112 from the imaging unit 110 at the input port. The hardware image splitter 130 may be operatively connected to the imaging unit 110 via a coaxial, twisted pair, fiber optic, or other suitable cable types for carrying image data. The hardware image splitter 130 may then split the input image data 112 (e.g., by copying of the stream of image data onto two or more transmission channels without the use of software) such that a copy of the image data 112 is presented for

output via the plurality of output ports. The hardware image splitter 130 may be compatible with one or more bus protocols, including Universal Serial Bus (USB) 3.2, USB 4.0, 10 Gigabit Ethernet, CoaXpress (CXP-12), CameraLink (Full), CameraLink (Deca), CameraLink HS (C3), or any other suitable bus protocol. The selection of a suitable bus protocol is important and non-trivial as it needs to be able to handle splitting data at high speed. Some non-limiting examples of a hardware image splitter 130 include the BitFlow CXP Replicator and the Vivid Engineering CLV-412.

[0030] As illustrated, the AVI environment 100 may include any number of image analysis hardware units (e.g., one, two, three, five, and so on). The one or more image analysis hardware units 140A-140N may include one or more processors configured to execute respective visual inspection analyses. The one or more processors may be an industrial PC that includes one or more processors that execute software instructions stored in a memory unit to perform some or all of the functions of the image analysis hardware units 140A-140N as described herein. The processors may be application-specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), graphics processing units (GPUs), computer processing units (CPUs), etc. Accordingly, the functionality of the image analysis hardware units 140A-140N as described herein may instead be implemented, in part or in whole, by such hardware. The memory units may include one or more physical memory devices or units containing volatile and/or non-volatile memory. Any suitable memory type or types may be used, such as read-only memory (ROM), solid-state drives (SSDs), hard disk drives (HDDs), and so on.

[0031] The image analysis hardware units 140A-140N may include one or more network interfaces. The network interfaces may include any suitable number and type of network interfaces, such as wired (*e.g.*, Ethernet, serial) and wireless (*e.g.*, WiFi, Bluetooth, etc.) and facilitate bidirectional communication over one or more networks. The network interfaces may also include one or more bus interfaces. The bus interfaces may include any suitable number and type of bus interfaces, such as USB, 10 Gigabit Ethernet, CoaXpress, and CameraLink.

[0032] The image analysis hardware units 140A-140N may receive the image data 112 from the output ports of the hardware image splitter 130. The image analysis hardware units 140A-140N may be connected to the hardware image splitter 130 via a bus interface connected to a coaxial, twisted pair, fiber optic, or other suitable cable coupled to a respective output port of the hardware image splitter 130. The image analysis hardware units 140A-140N may then analyze the image data 112 using visual inspection techniques, such as air gap measurements, plunger depth measurements, agitation, spinning, spin-stop, illumination with polarized light, LEDs, lasers, fluorescent lamps, incandescent lamps, or flash lamps, to detect potential product defects. In the pharmaceutical product context, some non-limiting examples of product defects include defective seals, plungers, or pistons, low fill, high fill, collapsed cake, meltback, liquefaction, stains, color variations, surface defects, dimensional variance, missing parts, misaligned parts, foreign particles, discoloration, and cracked or scratched vials.

[0033] In some embodiments, each of the image analysis hardware units 140A-140N may analyze the image data 112 using respective visual inspection techniques. For example, the image analysis hardware unit 140A may be

configured to detect cracks and/or scratches in a product surface and an image analysis hardware unit 140B may be configured to detect foreign contaminants. In some embodiments, the image analysis techniques implement machine learning (ML) / artificial intelligence (AI) techniques, such as deep learning, to identify characteristics of products 120 indicative of a defect. It should be appreciated that the disclosed AI/ML techniques are one example class of visual inspection analyses that may be applied to the image data 112 to detect characteristics of the product 120. Other visual inspection analyses that do not involve AI/ML are also envisioned.

[0034] In ML/AI embodiments, the memory may store a ML classifier trained to accept the image data 112 as an input and output a label indicative of whether or not the image data 112 includes one or more defects. To this end, the ML classifier may be trained using historical image data of products that exhibit and do not exhibit the defect of interest. The image analysis hardware units 140A-140N may then store the data representative of the trained ML classifier in a memory thereat. In some embodiments, different image analysis hardware units 140A-140N may implement similar ML techniques but are trained using different sets of training data.

[0035] In one embodiment, the trained ML classifier is an AVI neural network. The AVI neural network may classify entire images (e.g., defect vs. no defect, or presence or absence of a particular type of defect such as a crimp bruise or crimp defect generally, etc.), detect objects in images (e.g., detect the position of foreign objects that are not bubbles within container images), or some combination thereof (e.g., one neural network classifying images, and another performing object detection). "Object detection" broadly refers to techniques that identify the particular location of an object (e.g., a particle, a fiber, etc.) within an image, and/or that identify the particular location of a feature of a larger object (e.g., a bruised crimp or seal, a crack or chip on a syringe or cartridge barrel, etc.), and can include, for example, techniques that perform segmentation of the container image or image portion (e.g., pixel-by-pixel classification), or techniques that identify objects and place bounding boxes (or other boundary shapes) around those objects. The AVI neural network may include deep learning software such as MVTec from HALCON®, Vidi® from Cognex®, Rekognition® from Amazon®, TensorFlow, PyTorch, and/or any other suitable off-the-shelf or customized deep learning software. The AVI neural network software may be built on top of one or more pre-trained networks, such as ResNet50 or VGGNet, for example, and/or one or more custom networks. It should be appreciated that the specific ML structures used to train the ML classifier may vary depending on the particular analysis being performed.

[0036] It should be appreciated that because different image analysis hardware units 140A-140N may store different ML classifiers, the AVI environment 100 is able to analyze in parallel the image data 112 to increase the types of visual inspection analyses that can be performed in the processing window for a given product 120. Additionally, because the visual inspection techniques can be performed in parallel using a single set of image data, an AVI environment 100 may be configured to have fewer imaging stations than a conventional AVI environment that captures a new set of image data for each analysis that needs to be performed. As a result, the AVI environment

100 may have reduced complexity and may have a smaller footprint than a conventional AVI environment. Additionally, the parallel analyses performed by the image analysis hardware units 140A-140N may enable operators of the AVI environment 100 to evaluate performance of new ML classifiers as compared to currently-implemented ML classifiers without impacting operation of the AVI processes implemented at the AVI environment 100.

[0037] Figure 4 depicts an exemplary timing diagram 400 of an AVI cycle, according to some aspects. The horizontal axis of the exemplary timing diagram 400 represents time in milliseconds.

[0038] Each bar illustrates the elapsed time of an image inspection or image saving process. For example, processes 410, 420, and 430 may correspond to image inspections performed by image analysis hardware units 140A-140N, and process 440 may correspond to image saving performed by image storage hardware unit 150. Each process 410-440 may last a different duration, such as 75msec for process 410, 30msec for process 420, 50msec for process 430, and 100msec for process 440. As illustrated, each process 410-440 occurs concurrently, beginning at time 0msec, for each AVI cycle. As such the total processing time may be equal to the greatest duration of the individual processes 410-440, *i.e.*, 100msec.

[0039] In contrast, if the processes 410-440 were performed consecutively, *e.g.*, by a single hardware unit, the total processing time may be the sum of the durations of the individual processes 410-440, *i.e.*, 255msec, which may exceed the processing window duration of the production line.

[0040] Returning to Figure 1, the image analysis hardware units 140A-140N may generate image processing metadata 142A-142N based on the outputs of the visual inspection analyses performed thereat. As one example, the image processing metadata 142A-142N may include any labels output from the visual inspection analyses and a location of the defect in the image data 112 (e.g., indications defining a bounding box, an identifier indicating a particular frame of image data, etc.). As another example, the image analysis hardware units 140A-140N may analyze the image data 112 for product identification indicators, such as barcodes or labels. In this example, the image processing metadata 142A-142N may include a product and/or lot identifier, an image analysis hardware unit identifier, a timestamp, a defect identifier, and/or any other suitable information. The image analysis hardware units 140A-140N may then transmit the respective image processing metadata 142A-142N to the image storage hardware unit 150 for further processing. Additionally, the image analysis hardware units 140A-140N may compile the outputs of the respective image inspection analyses into a batch data document to provide a compilation of the results of the visual inspection analyses.

[0041] The image storage hardware unit 150 may include one or more processors, memories, network interfaces, bus interfaces, and buffers. The one or more processors may be an industrial PC that executes software instructions stored in a memory unit to perform some or all of the functions of the image storage hardware unit 150 as described herein. Alternatively, some of the processors may be other types of processors (e.g., application-specific integrated

circuits (ASICs), field-programmable gate arrays (FPGAs), graphics processing units (GPUs), etc.), and some of the functionality of the image storage unit 150 as described herein may instead be implemented, in part or in whole, by such hardware. The memory units may include one or more physical memory devices or units containing volatile and/or non-volatile memory. Any suitable memory type or types may be used, such as read-only memory (ROM), solid-state drives (SSDs), hard disk drives (HDDs), and so on. The buffer may comprise one or more solid-state drives (*e.g.*, PCIe 4.0 or 5.0), redundant arrays of inexpensive disks (RAID), or other suitable storage.

[0042] The network interfaces of the image storage hardware unit 150 may include any suitable number and type of network interfaces, such as wired (*e.g.*, Ethernet, serial) and wireless (*e.g.*, WiFi, Bluetooth, etc.) and facilitate bidirectional communication over one or more networks. The network interfaces may also include one or more bus interfaces. The bus interfaces may include any suitable number and type of bus interfaces, such as USB, 10 Gigabit Ethernet, CoaXpress, and CameraLink.

[0043] The image storage hardware unit 150 may receive image data 112 from an output port of the hardware image splitter 130. In some embodiments, the image storage hardware unit 150 may receive image acquisition metadata 114 from the output port of the hardware image splitter 130. The image analysis hardware units 140A-140N may be connected to the hardware image splitter 130 via a bus interface connected to a coaxial, twisted pair, fiber optic, or other suitable cable. Additionally, the image storage hardware unit 150 may receive (i) image processing metadata 142A-142N from the image analysis hardware units 140A-140N (ii) image acquisition metadata 114 from the imaging units 110, and/or (iii) system metadata 172 over a network, as a local area network (LAN) or a wired connection (e.g., via a serial or Ethernet port). The image storage hardware unit 150 may store the image data 112 and the inspection metadata (e.g., the image processing metadata 142A-142N, the image acquisition metadata 114, and/or the system metadata 172) in the buffer. More particularly, the image storage hardware unit 150 may stream the bits received over the bus interface from the hardware image splitter 130 into the buffer. While the image storage hardware unit 150 is in the process of storing the image data 112, the image storage hardware unit 150 may receive the image processing metadata 142A-142N from the image analysis hardware units 140A-140N and/or may receive the system environment metadata 172.

[0044] Depending on the format of the image data 112, the image storage hardware unit 150 may associate the image acquisition metadata 114, the system environment metadata 172, and/or the image processing metadata 142A-142N in different ways. For example, the image data 112 may be stored in a graphics file, such as joint photographic experts group (JPEG), tagged image file format (TIFF), portable network graphics (PNG), RAW, or other suitable file formats. For some of the file types, such as PNG files, metadata may be stored in a header portion of the file. Accordingly, for these file types, the image storage hardware unit 150 may associate the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N with the image data 112 by updating the header of the image data file to include the image acquisition metadata 114, the

system metadata 172, and/or the image processing metadata 142A-142N. For file types that do not support headers, the image storage hardware unit 150 may associate the image data 112 with the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N by modifying a file name for the image data 112 to indicate the metadata. Additionally, or alternatively, the image storage hardware unit 150 may create a simple database at the image storage hardware unit 150. In these embodiments, the image storage hardware unit 150 may store the image acquisition metadata 114, the system environment metadata 172, and/or the image processing metadata 142A-142N in a particular location in the buffer or a separate data store and the database may indicate the location of the image data 112 and the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N.

[0045] Regardless of the particular manner in which the image data 112 is associated with the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N in the buffer, the image storage hardware unit 150 may transfer the image data 112 and the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N to long-term storage after the association is complete to facilitate offline analysis of the image data 112. In some embodiments, the image storage hardware unit 150 may only transfer image data associated with a defect to the long-term storage. In these embodiments, if the image data 112 is not associated with a defect, the image storage hardware unit may simply allow the buffer to be repopulated with a stream of image data associated with the next product 120. In some embodiments, the image storage hardware unit 150 may be configured to compress the image data 112 and the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N before transfer to the long-term storage for more efficient data transfer and/or storage.

[0046] In some embodiments, the long-term storage is the cloud storage 160. In these embodiments, the image storage hardware unit 150 transmits the image data 112 and the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N to the cloud storage 160 via one or more local area networks (LANs) or wide area networks (WANs), such as the Internet. The cloud storage 160 may comprise one or more public clouds, such as Microsoft Azure, Amazon Web Services (AWS), and Google Cloud Platform, or private clouds. In some embodiments, the image storage hardware unit 150 may include a transmission queue in which a copy of the image data 112 and the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N is stored while awaiting transmission such that the buffer can begin receiving the image data 112 associated with a new product 120.

[0047] Additionally, in some embodiments, the image storage hardware unit 150 may be coupled to a local long-term storage unit (not depicted), such as a solid-state drive (SSD). In these embodiments, the long-term storage may be a separate memory unit within the image storage hardware unit 150 or an external storage unit coupled via a wired connection (e.g., via a data port). In these embodiments, a technician local to the AVI environment 100 is able

to manually review whether the products 120 have defects based on data maintained on-site and/or verify whether the visual inspection techniques implemented by the visual inspection hardware units 140A-140N are exhibiting false positives and/or negatives. This may enable faster correction of a source of defects while still maintaining the data security of locally-maintained data.

[0048] Turning to Figure 2, depicted is an exemplary image information database table 200, according to some aspects. The exemplary image information database table 200 may associate the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N with the image data 112. The exemplary image information database table 200 may be maintained on the image storage hardware unit 150 and/or the cloud storage 160. The image information database table can further include additional metadata such as a defect dimension or an object attribute. For example, the estimated size of a particle or the distance between two straight edges (in a plunger depth measurement).

[0049] The exemplary image information database table 200 may comprise information obtained from the image data 112 and the image acquisition metadata 114, the system metadata 172, and/or the image processing metadata 142A-142N. The exemplary image information database table 200 may include a plurality of fields 210-290 and a plurality of records. In the illustrated example, each row corresponds to one set of image data 112.

[0050] The image set number field 210 may include an image set identifier that identifies a set of image data associated with a particular product 120. The filename field 220 may include the filename of the image set. The date field 230 may include the month, day, and/or year that an image was generated or a product 120 was inspected. The time field 240 may include the time, *i.e.*, hour, minute, seconds, and/or milliseconds, that the image was generated or the product 120 was inspected. The shutter speed field 250 may include the shutter speed used to capture the image data 112. The product serial number field 260 may include a serial number that was read from the product 120, such as from a barcode or printed characters. The lot ID field 270 may include a lot identifier that was read from the product 120. The particles field 280 may indicate the presence and/or size of particles. The plunger depth field 290 indicate the depth of a syringe plunger. Additional or alternative fields may include any relevant Boolean, text, numerical, or other type of data about the product 120 and/or defect associated therewith. As some examples, additional fields may indicate a location at the image storage hardware unit 150 at which the image data 112 and the inspection metadata 142 are stored for each set of image data.

[0051] Figure 3 depicts an exemplary graphics file 300 with embedded image acquisition metadata 114 and image processing metadata 142A-142N, according to some aspects. In one aspect, the exemplary graphics file 300 may represents a PNG format file. As illustrated, the exemplary graphics file 300 may comprise a header 310, one or more metadata chunks 320-380, and a body 390.

[0052] The header 310 may include an image header (IHDR) and palette table (PLTE). The metadata chunks 320-380 may include the image acquisition metadata 114, the system metadata 172, and/or image processing metadata 142A-142N. The metadata chunks 320-380 may comprise ancillary chunks, such as textual data (tEXt), compressed textual data (zTXt), and/or international textual data (iTXt). The metadata chunks 320-380 may include a keyword, such as image no., date, time, shutter speed, product serial no., lot ID, and/or defects, and corresponding text string. The text string may comprise Boolean, text, or numerical data. The body 390 may include image data chunks (IDAT) derived from the image data 112 and an image trailer (IEND).

[0053] Figure 5 depicts a flow diagram of exemplary AVI method 500. One or more steps of the exemplary AVI method 500 may be implemented as a set of instructions stored on a computer-readable memory and executable on one or more processors. The exemplary AVI method 500 may be implemented via one or more systems, such as the imaging unit 110, hardware image splitter 130, image analysis hardware units 140A-140N, image storage hardware unit 150, and/or cloud storage 160.

[0054] In one embodiment, the exemplary AVI method 500 may include at block 510 generating image data of a product. The imaging unit (such as the imaging unit 110 and/or a digital camera) may generate the image data (such as the image data 112 and/or a set of image data referenced by the filename field 220). The imaging unit may include two or more cameras or image sensors oriented to capture images of the product at different orientations. The image data may include the images generated by the two or more cameras or image sensors. The imaging unit may generate outputs of at least 150 MB of image data per product (such as the products 120). The imaging unit may generate image acquisition metadata 114 providing information about the image data 112.

[0055] In one embodiment, the exemplary AVI method 500 may include at block 520 receiving the image data at a first image analysis hardware unit (such as one of the image analysis hardware units 140A-140N). The first image analysis hardware unit may receive the image data from a hardware image splitter (such as the hardware image splitter 130).

[0056] In one embodiment, the exemplary AVI method 500 may include receiving the image data at a second image analysis hardware unit (such as a second one of the image analysis hardware units 140A-140N. Embodiments may include one or more additional image analysis hardware units. The second image analysis hardware unit may receive the image data from the hardware image splitter.

[0057] In one embodiment, the exemplary AVI method 500 may include at block 530 analyzing the image data for product defects by applying a first visual inspection analysis and/or second visual inspection analysis. The first image analysis hardware unit may apply the first visual inspection analysis. The second image analysis hardware unit may apply the second visual inspection analysis. The first visual inspection analysis may include performing a particle detection analysis or a crack detection analysis. The second visual inspection analysis may include using a

trained AI algorithm, wherein the trained AI algorithm includes an ML classifier or an AVI neural network. Embodiments may include any number of additional image analysis hardware units performing respective visual inspection analyses.

[0058] In one embodiment, the exemplary AVI method 500 may include at block 540 generating first image processing metadata, such as one of the image processing metadata 142A-142N, based on the first visual inspection analysis. The exemplary AVI method 500 may include generating second image processing metadata, such as another one of the image processing metadata 142A-142N, based on the second visual inspection analysis. The first image analysis hardware unit may generate the first image processing metadata. The second image analysis hardware unit may generate the second image processing metadata. Embodiments may any number of additional image processing metadata generated by additional image analysis hardware units.

[0059] In one embodiment, the exemplary AVI method 500 may include at block 550 receiving the image data at an image storage hardware unit (such as the image storage hardware unit 150). The image storage hardware unit may receive the image data from the hardware image splitter. In one embodiment, the exemplary AVI method 500 may include receiving the image acquisition metadata 114 at the image storage hardware unit. The image storage hardware unit may receive the image acquisition metadata 114 from the hardware image splitter. In one embodiment, the exemplary AVI method 500 may include receiving the system metadata 172 at the image storage hardware unit. The image storage hardware unit may receive the system metadata 172 from the system controller unit 170.

[0060] In one embodiment, the exemplary AVI method 500 may include at block 560 storing the image data into a buffer. The image storage hardware unit may store the image data. The image data may be stored in PNG file format. In one embodiment, the exemplary AVI method 500 may include storing the image acquisition metadata 114 and/or system environment metadata 172 into the buffer. The image storage hardware unit may store the image acquisition metadata 114 and/or system environment metadata 172.

[0061] In one embodiment, the exemplary AVI method 500 may include at block 570 receiving the first image processing metadata and/or the second image processing metadata. The image storage hardware unit may receive the first image processing metadata from the first image analysis hardware unit and the second image processing metadata from the second image analysis hardware unit.

[0062] In one embodiment, the exemplary AVI method 500 may include at block 580 associating inspection metadata with the image data. The inspection metadata may include the first image processing metadata, the second image processing metadata, the image acquisition metadata, and/or the system control metadata. The image storage hardware unit may associate the inspection metadata with the image data in the buffer. The inspection metadata may be associated by embedding the inspection metadata into a PNG file. The inspection

metadata may be associated by associating the inspection metadata with an identifier of the image data in a database.

[0063] In one embodiment, the imaging unit may be coupled to the hardware image splitter via a first bus, the first image analysis hardware unit may be coupled to the hardware image splitter via a second bus, and the image storage hardware unit may be coupled to the hardware image splitter via a third bus. The first bus, second bus, and third bus may comprise CameraLink buses, USB buses, Gigabit Ethernet buses, or CoaXPress buses.

[0064] In one embodiment, the throughput of the exemplary AVI method 500 is at least 300 products per minute.

[0065] In one embodiment, the exemplary AVI method 500 includes transmitting the image data and associated metadata to a cloud storage system. The image storage hardware image may transmit the image data and associated metadata to a cloud storage (such as the cloud storage 160).

[0066] Additional considerations pertaining to this disclosure will now be addressed.

[0067] Some of the figures described herein illustrate example block diagrams having one or more functional components. It will be understood that such block diagrams are for illustrative purposes and the devices described and shown may have additional, fewer, or alternate components than those illustrated. Additionally, in various embodiments, the components (as well as the functionality provided by the respective components) may be associated with or otherwise integrated as part of any suitable components.

[0068] Embodiments of the disclosure relate to a non-transitory computer-readable storage medium having computer code thereon for performing various computer-implemented operations. The term "computer-readable storage medium" is used herein to include any medium that is capable of storing or encoding a sequence of instructions or computer codes for performing the operations, methodologies, and techniques described herein. The media and computer code may be those specially designed and constructed for the purposes of the embodiments of the disclosure, or they may be of the kind well known and available to those having skill in the computer software arts. Examples of computer-readable storage media include, but are not limited to: magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROMs and holographic devices; magneto-optical media such as optical disks; and hardware devices that are specially configured to store and execute program code, such as ASICs, programmable logic devices ("PLDs"), and ROM and RAM devices.

[0069] Examples of computer code include machine code, such as produced by a compiler, and files containing higher-level code that are executed by a computer using an interpreter or a compiler. For example, an embodiment of the disclosure may be implemented using Java, C++, or other object-oriented programming language and development tools. Additional examples of computer code include encrypted code and compressed code. Moreover, an embodiment of the disclosure may be downloaded as a computer program product, which may be

transferred from a remote computer (e.g., a server computer) to a requesting computer (e.g., a client computer or a different server computer) via a transmission channel. Another embodiment of the disclosure may be implemented in hardwired circuitry in place of, or in combination with, machine-executable software instructions.

[0070] As used herein, the singular terms "a," "an," and "the" may include plural referents, unless the context clearly dictates otherwise.

[0071] As used herein, the terms "approximately," "substantially," "substantial" and "about" are used to describe and account for small variations. When used in conjunction with an event or circumstance, the terms can refer to instances in which the event or circumstance occurs precisely as well as instances in which the event or circumstance occurs to a close approximation. For example, when used in conjunction with a numerical value, the terms can refer to a range of variation less than or equal to ±10% of that numerical value, such as less than or equal to ±5%, less than or equal to ±4%, less than or equal to ±3%, less than or equal to ±2%, less than or equal to ±1%, less than or equal to ±0.5%, less than or equal to ±0.1%, or less than or equal to ±0.05%. For example, two numerical values can be deemed to be "substantially" the same if a difference between the values is less than or equal to ±10% of an average of the values, such as less than or equal to ±5%, less than or equal to ±4%, less than or equal to ±3%, less than or equal to ±2%, less than or equal to ±1%, less than or equal to ±0.5%, less than or equal to ±0.1%, or less than or equal to ±0.05%.

[0072] Additionally, amounts, ratios, and other numerical values are sometimes presented herein in a range format. It is to be understood that such range format is used for convenience and brevity and should be understood flexibly to include numerical values explicitly specified as limits of a range, but also to include all individual numerical values or sub-ranges encompassed within that range as if each numerical value and sub-range is explicitly specified.

[0073] While the present disclosure has been described and illustrated with reference to specific embodiments thereof, these descriptions and illustrations do not limit the present disclosure. It should be understood by those skilled in the art that various changes may be made and equivalents may be substituted without departing from the true spirit and scope of the present disclosure as defined by the appended claims. The illustrations are not necessarily drawn to scale. There may be distinctions between the artistic renditions in the present disclosure and the actual apparatus due to manufacturing processes, tolerances and/or other reasons. There may be other embodiments of the present disclosure which are not specifically illustrated. The specification (other than the claims) and drawings are to be regarded as illustrative rather than restrictive. Modifications may be made to adapt a particular situation, material, composition of matter, technique, or process to the objective, spirit and scope of the present disclosure. All such modifications are intended to be within the scope of the claims appended hereto. While the techniques disclosed herein have been described with reference to particular operations performed in a particular order, it will be understood that these operations may be combined, sub-divided, or re-ordered to form an equivalent

technique without departing from the teachings of the present disclosure. Accordingly, unless specifically indicated herein, the order and grouping of the operations are not limitations of the present disclosure.

What is claimed is:

1. An automated visual inspection (AVI) system comprising:

an imaging unit for generating image data of a product;

a hardware image splitter comprising an input port, a first output port, and a second output port, wherein the imaging unit is coupled to the input port;

a first image analysis hardware unit, coupled to the first output port, comprising one or more processors and a memory storing computer-readable instructions, wherein the instructions cause the first image analysis hardware unit to:

receive the image data from the imaging unit via the hardware image splitter,

analyze the image data for product defects by applying a first visual inspection analysis, and

generate first image processing metadata based on the first visual inspection analysis; and

an image storage hardware unit, coupled to the second output port, comprising one or more processors, a buffer, and a memory storing computer-readable instructions, wherein the instructions cause the image storage hardware unit to:

store the image data from the hardware image splitter into the buffer,

receive the first image processing metadata, and

associate inspection metadata with the image data, wherein the inspection metadata includes the first image processing metadata.

2. The AVI system of claim 1, wherein:

the imaging unit includes two or more cameras or image sensors oriented to capture images of the product at different orientations, and

the image data includes the images generated by the two or more cameras or image sensors.

3. The AVI system of claims 1 or 2, wherein the imaging unit is coupled to the hardware image splitter via a first bus, the first image analysis hardware unit is coupled to the hardware image splitter via a second bus, the image storage hardware unit is coupled to the hardware image splitter via a third bus.

4. The AVI system of claim 3, wherein the first bus, second bus, and third bus comprise CameraLink buses, universal serial bus (USB) buses, Gigabit Ethernet buses, or CoaXPress buses.

5. The AVI system of any one of claims 1-4, wherein storing the image data from the hardware image splitter into the buffer comprises saving the image data into a portable network graphics (PNG) file, and associating the inspection metadata with the image data comprises embedding the inspection metadata into the PNG file.

6. The AVI system of any one of claims 1-4, wherein associating the inspection metadata with the image data comprises associating the inspection metadata with an identifier of the image data in a database.

7. The AVI system of any one of claims 1-6, wherein a throughput of the AVI system is at least 300 products per minute.

8. The AVI system of any one of claims 1-7, wherein the imaging unit generates outputs of at least 150MB of image data per product.

9. The AVI system of any one of claims 1-8, wherein the first visual inspection analysis includes a particle detection analysis, a crack detection analysis, or a scratch detection analysis.

10. The AVI system of any one of claims 1-9, wherein the hardware image splitter further comprises a third output port, the AVI system further comprising:

a second image analysis hardware unit, coupled to the third output port, comprising one or more processors and a memory storing computer-readable instructions, wherein the instructions cause the second image analysis hardware unit to:

receive the image data from the imaging unit via the hardware image splitter,

analyze the image data for product defects by applying a second visual inspection analysis, and

generate second image processing metadata based on the second visual inspection analysis,

wherein the image storage hardware unit further receives the second image processing metadata and the inspection metadata includes the second image processing metadata.

11. The AVI system of claim 10, wherein the second visual inspection analysis includes applying a machine learning algorithm trained using image data of products having a product defect.

12. The AVI system of any one of claims 1-11, wherein:

the imaging unit is configured to generate image acquisition metadata, and

the inspection metadata includes the image acquisition metadata.

13. The AVI system of any one of claims 1-12, further comprising:

a system controller for (i) controlling one or more components of the AVI system in accordance with one or more control parameters, and (ii) generating system metadata based on the one or more control parameters,

wherein the inspection metadata includes the system metadata.

14. The AVI system of any one of claims 1-13, wherein the instructions further cause the image storage hardware unit to:

transmit the image data and the inspection metadata to a cloud storage system.

15. An automated visual inspection (AVI)-based method comprising:

generating, via an imaging unit, image data of a product;

receiving, at a first image analysis hardware unit, the image data from the imaging unit via a hardware image splitter;

analyzing, using the first image analysis hardware unit, the image data for product defects by applying a first visual inspection analysis;

generating, using the first image analysis hardware unit, first image processing metadata based on the first visual inspection analysis;

receiving, at an image storage hardware unit, the image data from the hardware image splitter;

storing, using the image storage hardware unit, the image data into a buffer;

receiving, at an image storage hardware unit, the first image processing metadata; and

associating, using the image storage hardware unit, inspection metadata with the image data, wherein the inspection metadata includes the first image processing metadata.

16. The AVI-based method of claim 15, wherein:

the imaging unit includes two or more cameras or image sensors oriented to capture images of the product at different orientations, and

the image data includes the images generated by the two or more cameras or image sensors.

17. The AVI-based method of claims 15 or 16, wherein the imaging unit is coupled to the hardware image splitter via a first bus, the first image analysis hardware unit is coupled to the hardware image splitter via a second bus, the image storage hardware unit is coupled to the hardware image splitter via a third bus.

18. The AVI-based method of claim 17, wherein the first bus, second bus, and third bus comprise CameraLink buses, universal serial bus (USB) buses, Gigabit Ethernet buses, or CoaXPress buses.

19. The AVI-based method of any one of claims 15-18, wherein storing the image data from the hardware image splitter into the buffer comprises saving the image data into a portable network graphics (PNG) file, and associating the inspection metadata with the image data comprises embedding the inspection metadata into the PNG file.

20. The AVI-based method of any one of claims 15-18, wherein associating the inspection metadata with the image data comprises associating the inspection metadata with an identifier of the image data in a database.

21. The AVI-based method of any one of claims 15-20, wherein a throughput of the AVI-based method is at least 300 products per minute.

22. The AVI-based method of any one of claims 15-21, wherein the imaging unit generates an output of at least 150MB of image data per product.

23. The AVI-based method of any one of claims 15-22, wherein the first visual inspection analysis includes performing a particle detection analysis, a crack detection analysis, or a scratch detection analysis.

24. The AVI-based method of any one of claims 15-23, further comprising:

receiving, at a second image analysis hardware unit, the image data from the hardware image splitter,

analyzing, using the second image analysis hardware unit, the image data for product defects by applying a second visual inspection analysis,

generating, using the second image analysis hardware unit, second image processing metadata based on the second visual inspection analysis, and

receiving, at the image storage hardware unit, the second image processing metadata and the inspection metadata includes the second image processing metadata.

25. The AVI-based method of claim 24, wherein the second visual inspection analysis includes applying a machine learning algorithm trained using image data of products having a product defect.

26. The AVI-based method of any one of claims 15-25:

wherein generating the image data further comprises generating image acquisition metadata, and the inspection metadata includes the image acquisition metadata.

27. The AVI-based method of any one of claims 15-26, further comprising:

receiving, by the image storage unit and from a system controller, system metadata, wherein the inspection metadata includes the system metadata.

28. The AVI-based method of any one of claims 15-27, further comprising:

transmitting the image data and the inspection metadata to a cloud storage system.