SYSTEM AND METHOD FOR FACILITATING COMPLIANCE AND BEHAVIORAL ACTIVITY VIA SIGNALS DRIVEN BY ARTIFICIAL INTELLIGENCE

CROSS REFERENCE TO RELATED APPLICATIONS

[1] The present application claims priority to and the benefit of Singapore Patent Application No. 10202250173Q, filed on June 15, 2022, the entirety of which is hereby incorporated by reference in its entirety.

FIELD OF THE INVENTION

[2] The present application relates to analytics technologies, behavioral compliance technologies, artificial intelligence technologies, machine learning technologies, cloud-computing technologies, data analysis technologies, and, more particularly, to a system and method for facilitating compliance and behavioral activity via signals driven by artificial intelligence.

BACKGROUND

[3] In today's society, it has become increasingly desirable and important to be able to motivate individuals to follow regimens, care plans, and goals in a manner that is both effective and lasting. Such desirability and importance are especially evident when it comes to motivating individuals with respect to maintaining compliance with medical regimens, weight loss programs, personal goals for physical and/or mental achievement, training objectives, removing bad habits, among other plans and goals. Facilitating compliance and shaping behavioral activity in accordance with regimens, care plans, and/or other goals often leads to desirable outcomes. To attempt to achieve such desirable outcomes, individuals often rely on using written or digital instructions that they may or may not follow, personal coaches that may utilize various techniques to motivate individuals, personally setting reminders in calendars or various types of software applications, attempting to pair up with other individuals with common goals and interests, or self-motivation. While such forms of facilitating compliance may achieve some level of compliance and/or behavioral change, they are often inconsistent,

difficult to comply with, do not stimulate individuals into action, or are simply ineffective at achieving desirable outcomes.

[4] Based on at least the foregoing, there remains room for substantial enhancements to existing technologies and processes and for the development of new technologies and processes to facilitate compliance and stimulate behavioral activity. While currently existing technologies provide for various benefits, such technologies still come with various drawbacks and inefficiencies. For example, it is relatively easy for an individual to ignore doctor's instructions, medical regimens, reminders in calendars, and even personal coaches. Even if an individual does initially comply and/or change their behaviors, the individual may only do so for a short period of time before losing focus, losing motivation, or giving up. When individuals lose focus, motivation, and/or give up, not only may they relapse to their original starting point, but they may regress even further because they may believe that their goals cannot be achieved. Furthermore, while currently existing techniques and processes may have short-term effectiveness on a case-by-case basis, existing technologies often fail to have a lasting effect on changing behavior over the long term. Moreover, existing technologies fail to take advantage of artificial intelligence technologies that may assist in adapting to changing conditions that may affect individuals. Based on the foregoing, current technologies may be improved and enhanced to provide for more effective techniques for motivating individuals, ensuring compliance with regimens, more effective intervention and/or response processes, higher quality predictive capabilities, and greater participation in various types of programs. Such enhancements and improvements to methodologies and technologies may provide for enhanced compliance, lasting behavioral change, and significantly improved outcomes.

SUMMARY

[5] A system and accompanying methods for facilitating compliance and behavioral activity via signals (e.g., participant's responses to cues) driven by artificial intelligence are disclosed. In certain embodiments, the system and methods involve utilizing data and artificial intelligence models to effectively obtain and analyze data associated with individuals (i.e., participants of the system) from a variety of data sources; predict, compute, and update markers

associated with the individuals based on the obtained data; generate insights relating to an individual's past and current markers; generate recommendations for the next best action or activity for an individual to perform to advance towards a goal; generate a ranked list of cues for the individual to interact with to motivate the individual to perform the action or activity; create segments to target cues towards specific sub-populations within a set of individuals; monitor the progress of each individual in terms of advancing towards the individual's goal; obtain telemetry data from systems and devices that have data associated with interactions with the cues and/or individuals' progress towards goals; and utilize the telemetry data to update the system and models utilized to facilitate the operative functionality of the system.

[6] In certain embodiments, the system and methods may include determining goals for individuals and capturing content and/or data from a variety of different data sources, such as devices and/or systems that may have data associated with an individual. Such content and/or data may include, but is not limited to, demographic data, psychographic data, health behavior data (e.g., physical activity, nutrition, sleep, and/or other health-related data), program participation data, health indicator data (e.g., health screening or biometric data), along with other types of data that may be associated with an individual. The captured content and/or data may be loaded into data models and artificial intelligence models that have been trained to recognize patterns, behaviors, objects, activities, individuals, and/or other items of interest. Such artificial intelligence models may be trained to recognize the patterns, behaviors, objects, activities, individuals, and/or other items of interest based on analyzing other content and/or data that have been fed into the models on previous occasions. In certain embodiments, the effectiveness and detection capability of the artificial intelligence models may be enhanced as the models receive additional content and/or data over time. The captured content and/or data may be compared to the content and/or data used to train the models and/or to deductions, reasoning, intelligence, correlations, outputs, analyses, and/or other information that the artificial intelligence model(s) learned based on the content and/or data used to train the models.

[7] Once the content and/or data are captured from the data sources, microbots

powered via artificial intelligence models may be utilized to compute, predict, and update markers for each individual as new data arrives at the system. The microbots perform the foregoing functions to facilitate generation of the most relevant cues for individuals to interact with. The markers and data utilized to compute the markers may be stored in databases to provide insights into each individual's current and past sets of markers. The databases may also be utilized to record the specific pathways that individuals have taken towards achieving goals, as their markers change over time. A recommender system of the system may utilize neural networks (e.g., graph neural networks) and sequence-based recommender functionality to generate recommendations for actions for each individual that, when performed, would advance the individual towards the individual's goal. Additionally, the recommender system may determine a list of cues that correspond with the actions to be performed for each individual. In certain embodiments, the list of cues may be a ranked list of cues with the top ranked cue having the highest probability of being opened by an individual and/or the highest probability of motivating the individual to perform the action in response to interacting with the cue.

[8] In certain embodiments, the cues may include audio, visual, virtual reality, augmented reality, text, electronic messages, and/or any other type of perceivable content, which may be provided to various delivery channels including mobile applications, email, simple messaging service, calls, and/or other delivery channels to deliver the cues to each of the individuals. In certain embodiments, a companion application of the system may allow users to author cues to engage with each of the individuals and may also create segments to target cues at specific sub-populations of individuals. Mobile applications may be provided that monitor the daily progress of each of the individuals, such as whether daily health behaviors have been performed and progress towards goals have been made. Similarly, regression and/or stagnation with respect to goals may also be monitored. The various applications that are utilized to provide the cues to the individuals may generate telemetry data, which may be provided to the system as individuals interact with the cues (or do not interact with the cues). The telemetry data may include information relating to the interaction with cues, information relating to changes of knowledge graphs used with the system, outputs that

may be utilized to train models of the system, any other information, or a combination thereof. In certain embodiments, the telemetry data may be utilized to update the models utilized by the system so that the system becomes more effectively at determining actions to perform and cues to motivate individuals to perform the actions over time.

[9] In certain embodiments, a system for facilitating compliance and behavioral activity via signals driven by artificial intelligence is provided. The system may include a memory that stores instructions and a processor that executes the instructions to perform various operations of the system. The system may perform an operation that includes capturing data from a plurality of data sources including information associated with individuals and determining goals for the individuals. Additionally, the system may perform an operation that includes determining markers associated with the individuals based on the data, such as by utilizing artificial intelligence models. Based on the markers, segments, topics, and/or other information, the system may determine, such as by utilizing the artificial intelligence models, actions for the individuals to perform and one or more cues for motivating the individuals to perform the actions. In certain embodiments, the system may include performing an operation that includes ranking the cues based on a probability of an individual opening the cue and/or performing an action in response to the cue that may be utilized to advance an individual towards his or her goal. The system may include performing an operation that includes providing the cues to the individuals for interaction. In certain embodiments, an individual's interactions (or lack of interactions) with the cues may be monitored and data relating thereto may be fed back into the system to update the artificial intelligence and data models utilized by the system to determine the actions and cues as the system operates over time. In certain embodiments, the system may monitor positive interactions, lack of interactions, negative interactions, negative behaviors conducted in response to cues, positive behaviors conducted in response to cues, indifferent behaviors conducted in response to cues, or a combination thereof.

[10] In certain embodiments, a method for facilitating compliance and behavioral activity via signals driven by artificial intelligence is disclosed. The method may include a memory that stores instructions and a processor that executes the instructions to perform the

functionality of the method. In particular, the method may include capturing data from a plurality of data sources including information associated with individuals and determining goals for the individuals. Additionally, the method may include determining markers associated with the individuals based on the data, such as by utilizing artificial intelligence models. Based on the markers, segments, topics, and/or other information, the method may include determining, such as by utilizing the artificial intelligence models, actions for the individuals to perform and one or more cues for motivating the individuals to perform the actions. In certain embodiments, the method may include ranking the cues based on a probability of an individual opening the cue and/or performing an action in response to the cue that may be utilized to advance an individual towards his or her goal. The method may include providing the cues to the individuals for interaction. In certain embodiments, an individual's interactions (or lack of interactions) with the cues may be monitored and data relating thereto may be fed back into a system performing the method to update the artificial intelligence and data models utilized by the system to determine the actions and cues as the method operates over time.

[11] In certain embodiments, a computer-readable device comprising instructions, which, when loaded and executed by a processor cause the processor to perform operations, the operations comprising: capturing data from a plurality of data sources including information associated with individuals and determining goals for the individuals; determining markers associated with the individuals based on the data, such as by utilizing artificial intelligence models; based on the markers, segments, topics, and/or other information, determining, such as by utilizing the artificial intelligence models, actions for the individuals to perform and one or more cues for motivating the individuals to perform the actions; ranking the cues based on a probability of an individual opening the cue and/or performing an action in response to the cue that may be utilized to advance an individual towards his or her goal; providing the cues to the individuals for interaction; monitoring an individual's interactions (or lack of interactions) with the cues and feeding data relating thereto into a system to update the artificial intelligence and data models utilized by the system to determine the actions and cues as the system operates over time.

[12] These and other features of the systems and methods for facilitating compliance and behavioral activity via signals driven by artificial intelligence are described in the following detailed description, drawings, and appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[13] Figure 1 is a schematic diagram of a system for facilitating compliance and behavioral activity via signals driven by artificial intelligence according to an embodiment of the present disclosure.

[14] Figure 2 is a schematic diagram of a system forming a part of or utilized in conjunction with the system of Figure 1, which illustrates an architectural overview of various components, systems, and functionality provided by the system according to an embodiment of the present disclosure.

[15] Figure 3 is a schematic diagram of a portion of the system of Figure 1, which processes and utilizes telemetry data to assist in the updating of models and graphs utilized by the system of Figure 1.

[16] Figure 4 is a schematic diagram illustrating customer and system development environments for facilitating model training and scoring on customer data according to an embodiment of the present disclosure.

[17] Figure 5 is a schematic diagram illustrating a cloud of the system interacting with customer clouds to exchange data that may be utilized to update models and functionality of the system of Figure 1.

[18] Figure 6 is a diagram illustrating various variables and data utilized by the system of Figure 1 to facilitate the generation of cues to facilitate compliance and behavioural activity by users.

[19] Figure 7 is a diagram illustrating the ranking of cues and batch-based transmission of the cues to applications for ultimate delivery to users for interaction.

[20] Figure 8 is a diagram illustrating on-demand ranking and delivery of cues to users for interaction.

[21] Figure 9 is a sample system architecture for providing and generating rankings of cues and relationships of the cues with variables utilized by the system of Figure 1.

[22] Figure 10 is an illustrative example depicting sequential information that may be utilized by graphical neural networks of the system of Figure 1 to rank cues and personalize pathways for users.

[23] Figure 11 illustrates algorithms that may be utilized for recommendation system supporting the functionality of Figure 1.

[24] Figure 12 illustrates exemplary candidate item generation for use with the system of Figure 1.

[25] Figure 13 illustrates further detail relating to candidate item generation.

[26] Figure 14 illustrates further detail relating to candidate item generation and issues relating thereto.

[27] Figure 15 illustrates further detail relating to candidate item generation and issues relating thereto.

[28] Figure 16 illustrates possible solutions for addressing potential issues associated with candidate item generation.

[29] Figure 17 illustrates an exemplary implementation of a recommendation system as a graph with deep learning.

[30] Figure 18 illustrates further details relating to an implementation of a recommendation system as a graph with deep learning.

[31] Figure 19 illustrates further details relating to an implementation of a recommendation system as a graph with deep learning.

[32] Figure 20 illustrates further details relating to an implementation of a recommendation system as a graph with deep learning.

[33] Figure 21 illustrates session-based recommendation capability of the system.

[34] Figure 22 illustrates various challenges associated with session-based recommendation systems.

[35] Figure 23 illustrates exemplary solutions for session-based recommendation systems for use with the system of Figure 1.

[36] Figure 24 illustrates the construction of session graphs for use with the system of Figure 1.

[37] Figure 25 illustrates learning item embeddings on session graphs for use with the system of Figure 1.

[38] Figure 26 illustrates generating session embeddings.

[39] Figure 27 illustrates making recommendations and model training for the system of Figure 1.

[40] Figure 28 is an exemplary system architecture illustrating the generation of cues based on various variables analysed by the system of Figure 1.

[41] Figure 29 is a schematic diagram illustrating a pipeline for use with the system of Figure 1 that is in a training mode for training models utilized by the system of Figure 1.

[42] Figure 30 is a schematic diagram illustrating a pipeline for use with the system of Figure 1 that is in an inference mode that includes generating a ranked list of cues.

[43] Figure 31 is a schematic diagram illustrating a pipeline for use with the system of Figure 1 that is in an inference mode that provides further detail relating to generating a ranked list of cues.

[44] Figure 32 is a schematic diagram illustrating a pipeline for use with the system of Figure 1 that is in an inference mode that provides further detail relating to responses to cues that are made by users interacting with the cues.

[45] Figure 33 illustrates an exemplary knowledge graph attention network that may be utilized with the system of Figure 1 to facilitate generation of recommendations for cues.

[46] Figure 34 illustrates further details relating to an exemplary knowledge graph attention network that may be utilized with the system of Figure 1 to facilitate generation of recommendations for cues.

[47] Figure 35 illustrates further details relating to a knowledge graph attention network that may be utilized with the system of Figure 1.

[48] Figure 36 illustrates utilizing a knowledge graph as an input to the system of Figure 1 and generating an output of recommended cues by utilizing the system of Figure 1.

[49] Figure 37 illustrates exemplary parameters for use with a knowledge graph attention network utilized in the system of Figure 1.

[50] Figure 38 illustrates a table containing exemplary experimental results based on training the knowledge graph attention network.

[51] Figure 39 is a schematic diagram illustrating the relationships between participants, cues, markers, segments, and topics in a knowledge graph utilized by the system of Figure 1.

[52] Figure 40 is a schematic diagram illustrating personalization of cue recommendations facilitated by a knowledge graph utilized by the system of Figure 1.

[53] Figure 41 is a schematic diagram illustrating personalization of cue recommendations facilitated by a knowledge graph utilized by the system of Figure 1.

[54] Figure 42 is a schematic diagram illustrating an ability of the system of Figure 1 to adapt to changes in user behavior.

[55] Figure 43 is a schematic diagram illustrating further detail relating to an ability of the system of Figure 1 to adapt to changes in user behavior.

[56] Figure 44 is a schematic diagram illustrating further detail relating to an ability of the system of Figure 1 to adapt to changes in user behavior.

[57] Figure 45 is a schematic diagram illustrating further detail relating to an ability of the system of Figure 1 to adapt to changes in user behavior.

[58] Figure 46 is a schematic diagram illustrating further detail relating to an ability of the system of Figure 1 to adapt to changes in user behavior.

[59] Figure 47 is a schematic diagram illustrating further detail relating to an ability of the system of Figure 1 to adapt to changes in user behavior.

[60] Figure 48 is a schematic diagram illustrating further detail relating to an ability of the system of Figure 1 to adapt to changes in user behavior.

[61] Figure 49 is a schematic diagram illustrating an ability of the system of Figure 1 to adapt the generation of cues to various areas of health for a user.

[62] Figure 50 is a schematic diagram illustrating the generation of on-demand cues by the system of Figure 1.

[63] Figure 51 illustrates an exemplary framework for use with the system of Figure 1.

[64] Figure 52 illustrates various exemplary design options for a framework for use with the system of Figure 1.

[65] Figure 53 illustrates a diagram depicting an ability to utilize learned embeddings for the same graph nodes from multiple graph deployments and combining the learnings to generate new versions of knowledge graphs utilized by the system of Figure 1.

[66] Figure 54 illustrates portions of a knowledge graph that are transferrable according to embodiments of the present disclosure.

[67] Figure 55 is a diagram illustrating the ability to create signatures for participants, cues, and segments according to embodiments of the present disclosure.

[68] Figure 56 illustrates a pair of examples for determining the number of possible participant signatures based on markers in the signatures.

[69] Figure 57 illustrates a diagram for reinforcement learning for use with the system of Figure 1.

[70] Figure 58 illustrates further detail relating to reinforcement learning for use with the system of Figure 1.

[71] Figure 59 illustrates further detail relating to reinforcement learning that involves utilizing data schema and collection processes across multiple deployments at various customers communicatively linked with the system of Figure 1.

[72] Figure 60 illustrates a diagram depicting adapting and adjusting a reinforcement training model as the model acquires new data.

[73] Figure 61 is a flow diagram illustrating a sample method for facilitating compliance and behavioral activity via signals driven by artificial intelligence according to an embodiment of the present disclosure.

[74] Figure 62 is a schematic diagram of a machine in the form of a computer system within which a set of instructions, when executed, may cause the machine to facilitate compliance and behavioral activity via signals driven by artificial intelligence.

DETAILED DESCRIPTION OF THE DRAWINGS

[75] In certain embodiments, a system 100 and accompanying methods (e.g., method 6100) for facilitating compliance and behavioral activity via signals (e.g., participants' responses to cues) driven by artificial intelligence are disclosed. The system 100 and methods utilize data and artificial intelligence models to obtain and analyze data from a variety of data sources. In certain embodiments, the data may be associated with individuals. The models of

the system 100 may be utilized to predict, compute, and update markers associated with the individuals based on the obtained data. In certain embodiments, the system 100 may generate insights relating to an individual's past and current markers and may generate recommendations for a next action or activity for an individual to perform to advance the individual towards a goal, such as, a health goal, work goal, activity goal, and/or other types of goals. Once the recommended actions are determined by the system 100, the system 100, in certain embodiments, may generate a ranked list of cues for the individual to interact with to motivate or nudge the individual to perform the recommended action or activity. In certain embodiments, the system 100 may create or determine segments to target cues towards specific sub-populations within a set of individuals. In certain embodiments, the cues may include text content, audio content, video content, augmented reality content, virtual reality content, haptic content, any type of content, or a combination thereof. Once generated and ranked, the cues may be pushed to each individual via various delivery channels, such as via third party applications in communication with the system 100, content delivery systems, mobile devices, and/or other systems and devices capable of delivering content to individuals. In certain embodiments, the content may be provided to human or robotic coaches that may utilize the cues in combination with coaching to interact with the participants.

[76] In certain embodiments, the system 100 may monitor the progress of each individual in terms of advancing towards the individual's goal and may determine which cues that the individual ultimately interacted with and led to action by the individual. In certain embodiments, the third-party applications, devices, and/or systems that provide the cues for consumption by the individuals may generate telemetry data associated with the cues (e.g., performance of the cues, which cues were opened/interacted with, etc.), associated with the individuals, associated with progression or regression with respect to a goal, and/or associated with any other information. The system 100 may obtain the telemetry data and may utilize the telemetry data to update the system 100 and models utilized to facilitate the operative functionality of the system 100 over time. Based on such updates, the system 100 and accompanying models may more effectively generate recommendations for activities to advance an individual towards a goal, generate more effective and meaningful cues for

individuals to interact with, and facilitate individual compliance and/or activity through the generation of increasingly effective recommended actions and cues over time.

[77] In certain embodiments, the system 100 and methods may determine goals for individuals and capture content and/or data from a variety of different data sources, such as devices and/or systems that may have data associated with an individual. For example, such content and/or data may include, but is not limited to, demographic data, psychographic data, health behavior data (e.g., physical activity, nutrition, sleep, and/or other health-related data), program participation data, health indicator data (e.g., health screening or biometric data), sensor data, body measurement data, sleep data, and other types of data that may be associated with an individual. The captured content and/or data may be loaded into data models and artificial intelligence models that have been trained to recognize patterns, behaviors, objects, activities, individuals, and/or other items of interest. In certain embodiments, the artificial intelligence models may be trained to recognize the patterns, behaviors, objects, activities, individuals, and/or other items of interest based on analyzing other content and/or data that have been fed into the models on previous occasions. In certain embodiments, the effectiveness and detection capability of the artificial intelligence models may be enhanced by the system 100 as the models receive additional content and/or data over time. The captured content and/or data may be compared by the system 100 to the content and/or data used to train the models and/or to deductions, reasoning, intelligence, correlations, outputs, analyses, and/or other information that the artificial intelligence model(s) learned based on the content and/or data used to train the models.

[78] Once the content and/or data are obtained from the data sources, microbots powered via artificial intelligence models may be utilized to compute, predict, and update markers for each individual as new data arrives at the system 100. In certain embodiments, the markers may comprise information that may be utilized to identify an individual and/or may comprise specific characteristics corresponding to the individual that may be determined from the content and/or data. The microbots perform the foregoing functions to facilitate generation of the most relevant cues for individuals to interact with. The markers and data utilized to

compute the markers may be stored in databases (e.g., database 155) to provide insights into each individual's current and past sets of markers. The databases may also be utilized to record the specific pathways that individuals have taken towards achieving goals as their markers change over time. In certain embodiments, a recommender subsystem of the system 100 may utilize neural networks (e.g., graph neural networks) and sequence-based recommender functionality to generate recommendations for actions for each individual which, when performed, would advance the individual towards the individual's goal. Additionally, in certain embodiments, the recommender system may determine a list of cues that correspond with the actions to be performed for each individual. In certain embodiments, the list of cues may be a ranked list of cues with the top ranked cue having the highest probability of being opened and/or interacted with by an individual and/or the highest probability of motivating the individual to perform the action based on interacting with the cue. In certain embodiments, a probability of a cue being opened, interacted with, and/or being able to motivate action may be based on a participant having a predilection to opening other cues having one or more characteristics that correlate with the cue, based on the cue having content correlating with one or more characteristics of the participant, based on the participant's favorable previous interaction with the cue, or a combination thereof. In certain embodiments, the probability may be measured between and including zero to one, and may be expressed in percentages, decimals, fractions, and/or other ways in which probabilities may be expressed.

[79] In certain embodiments, the cues may be provided to various delivery channels including mobile applications, email, simple messaging service, calls, and/or other delivery channels to deliver the cues to each of the individuals. In certain embodiments, a companion application of the system 100 or an application in communication with the system 100 may allow users to author cues to engage with each of the individuals and may also create segments to target cues at specific sub-populations of individuals. Applications may be provided that monitor the daily progress of each of the individuals, such as whether daily health behaviors have been performed and progress towards goals have been made. As indicated above, the various applications that are utilized to provide the cues to the individuals may generate telemetry data, which may be provided to the system 100 as individuals interact with the cues

(or do not interact with the cues). The telemetry data may include information relating to the interaction with cues, information relating to changes of knowledge graphs used with the system, outputs that may be utilized to train models of the system, any other information, or a combination thereof. In certain embodiments, the telemetry data may be utilized to update the models utilized by the system so that the system becomes more effective at determining actions to perform and cues to motivate individuals to perform the actions over time. As the models are updated over time, the generation of the cues, selection of the cues, ranking of the cues, and/or the content of the cues may be enhanced so that cues recommended by the system 100 in the future are enhanced and more effective in facilitating individual compliance and/or activity.

[80] As shown in Figure 1 and referring also to Figures 2-62, a system 100 for facilitating compliance and behavioral activity via signals driven by artificial intelligence are disclosed. In certain embodiments, the system 100 may be configured to support, but is not limited to supporting, compliance systems and services, behavioral activity systems and services, monitoring systems and services, alert systems and services, data analytics systems and services, data collation and processing systems and services, artificial intelligence services and systems, machine learning services and systems, content delivery services, cloud computing services, satellite services, telephone services, voice-over-internet protocol services (VoIP), software as a service (SaaS) applications, platform as a service (PaaS) applications, gaming applications and services, social media applications and services, operations management applications and services, productivity applications and services, mobile applications and services, and/or any other computing applications and services. In certain embodiments, the system 100 may include a first user 101, who may utilize a first user device 102 to access data, content, and services, or to perform a variety of other tasks and functions. As an example, the first user 101 may utilize first user device 102 to transmit signals to access various online services and content, such as those available on an internet, on other devices, and/or on various computing systems. As another example, the first user device 102 may be utilized to access an application, devices, and/or components of the system 100 that provide any or all of the operative functions of the system 100. In certain embodiments, the first user 101 may be a person, a robot, a humanoid, a program, a computer, any type of user, or a

combination thereof, that may seek to achieve a certain goal, such as comply with a regimen or program. In certain embodiments, the first user and/or any other user described herein may be any type of user, an individual, a participant of the system 100, an end user (e.g., an end user that receives and/or interacts with the cues), a user that manages various aspects of the system 100 (e.g., a user that authors cues and/or segments), or a combination thereof. In certain embodiments, any number and/or types of users may participate in the system 100. The first user device 102 may include a memory 103 that includes instructions, and a processor 104 that executes the instructions from the memory 103 to perform the various operations that are performed by the first user device 102. In certain embodiments, the processor 104 may be hardware, software, or a combination thereof. The first user device 102 may also include an interface 105 (e.g., screen, monitor, graphical user interface, etc.) that may enable the first user 101 to interact with various applications executing on the first user device 102 and to interact with the system 100. In certain embodiments, the first user device 102 may be and/or may include a computer, any type of sensor, a laptop, a set-top-box, a tablet device, a phablet, a server, a mobile device, a smartphone, a smart watch, and/or any other type of computing device. Illustratively, the first user device 102 is shown as a smartphone device in Figure 1. In certain embodiments, the first user device 102 may be utilized by the first user 101 to control and/or provide some or all of the operative functionality of the system 100.

[81] In addition to using first user device 102, the first user 101 may also utilize and/or have access to additional user devices. As with first user device 102, the first user 101 may utilize the additional user devices to transmit signals to access various online services and content. The additional user devices may include memories that include instructions, and processors that executes the instructions from the memories to perform the various operations that are performed by the additional user devices. In certain embodiments, the processors of the additional user devices may be hardware, software, or a combination thereof. The additional user devices may also include interfaces that may enable the first user 101 to interact with various applications executing on the additional user devices and to interact with the system 100. In certain embodiments, the first user device 102 and/or the additional user devices may be and/or may include a computer, any type of sensor, a laptop, a set-top-box, a

tablet device, a phablet, a server, a mobile device, a smartphone, a smart watch, and/or any other type of computing device, and/or any combination thereof. In certain embodiments, the sensors may include, but are not limited to, cameras, motion sensors, acoustic/audio sensors, pressure sensors, temperature sensors, light sensors, heart-rate sensors, blood pressure sensors, sweat detection sensors, breath-detection sensors, stress-detection sensors, body and/or vital sign measurement sensors, any type of health sensor, any type of sensors, or a combination thereof.

[82] The first user device 102 and/or additional user devices may belong to and/or form a communications network. In certain embodiments, the communications network may be a local, mesh, or other network that enables and/or facilitates various aspects of the functionality of the system 100. In certain embodiments, the communications network may be formed between the first user device 102 and additional user devices using any type of wireless or other protocol and/or technology. For example, user devices may communicate with one another in the communications network by utilizing any protocol and/or wireless technology, satellite, fiber, or any combination thereof. In certain embodiments, the communications network may be configured to communicatively link with and/or communicate with any other network of the system 100 and/or outside the system 100.

[83] In certain embodiments, the first user device 102 and additional user devices belonging to the communications network may share and exchange data with each other via the communications network. For example, the user devices may share information relating to the various components of the user devices, information associated with images and/or content accessed by a user of the user devices, information identifying the locations of the user devices, information indicating the types of sensors that are contained in and/or on the user devices, information identifying the applications being utilized on the user devices, information identifying how the user devices are being utilized by a user, information identifying user profiles for users of the user devices, information identifying device profiles for the user devices, information identifying the number of devices in the communications network, information identifying devices being added to or removed from the communications network,

any other information, or any combination thereof.

[84] In addition to the first user 101, the system 100 may also include a second user 110. The second user 110 may also be a person that may also seek to accomplish a goal or may have been prescribed a regimen to follow, such as by a physician, health professional, teacher, work colleague, or other individual. In certain embodiments, the second user device 111 may be utilized by the second user 110 to transmit signals to request various types of content, services, and data provided by and/or accessible by communications network 135 or any other network in the system 100. In further embodiments, the second user 110 may be a robot, a computer, a humanoid, an animal, any type of user, or any combination thereof. The second user device 111 may include a memory 112 that includes instructions, and a processor 113 that executes the instructions from the memory 112 to perform the various operations that are performed by the second user device 111. In certain embodiments, the processor 113 may be hardware, software, or a combination thereof. The second user device 111 may also include an interface 114 (e.g., screen, monitor, graphical user interface, etc.) that may enable the second user 110 to interact with various applications executing on the second user device 111 and, in certain embodiments, to interact with the system 100. In certain embodiments, the second user device 111 may be a computer, a laptop, a set-top-box, a tablet device, a phablet, a server, a mobile device, a smartphone, a smart watch, and/or any other type of computing device. Illustratively, the second user device 111 is shown as a mobile device in Figure 1. In certain embodiments, the second user device 111 may also include sensors, such as, but are not limited to, cameras, audio sensors, motion sensors, pressure sensors, temperature sensors, light sensors, heart-rate sensors, blood pressure sensors, oxygen sensors, sweat detection sensors, breath-detection sensors, stress-detection sensors, any type of health sensor, any type of sensors, or a combination thereof.

[85] In certain embodiments, the first user device 102, the additional user devices, and/or the second user device 111 may have any number of software applications and/or application services stored and/or accessible thereon. For example, the first user device 102, the additional user devices, and/or the second user device 111 may include applications for controlling and/or

accessing the operative features and functionality of the system 100, applications for controlling and/or accessing any device of the system 100, interactive social media applications, biometric applications, cloud-based applications, VoIP applications, other types of phone-based applications, product-ordering applications, business applications, e-commerce applications, media streaming applications, content-based applications, media-editing applications, database applications, gaming applications, internet-based applications, browser applications, mobile applications, service-based applications, productivity applications, video applications, music applications, social media applications, any other type of applications, any types of application services, or a combination thereof. In certain embodiments, the software applications may support the functionality provided by the system 100 and methods described in the present disclosure. In certain embodiments, the software applications and services may include one or more graphical user interfaces to enable the first and/or potentially second users 101, 110 to readily interact with the software applications. The software applications and services may also be utilized by the first and/or second users 101, 110 to interact with any device in the system 100, any network in the system 100, or any combination thereof. In certain embodiments, the first user device 102, the additional user devices, and/or the second user device 111 may include associated telephone numbers, device identities, or any other identifiers to uniquely identify the first user device 102, the additional user devices, and/or the second user device 111.

[86] The system 100 may also include a communications network 135. The communications network 135 may be under the control of a service provider, a manager of the system 100, the first user 101, any other designated user, a computer, another network, or a combination thereof. The communications network 135 of the system 100 may be configured to link each of the devices in the system 100 to one another. For example, the communications network 135 may be utilized by the first user device 102 to connect with other devices within or outside communications network 135. Additionally, the communications network 135 may be configured to transmit, generate, and receive any information and data traversing the system 100. In certain embodiments, the communications network 135 may include any number of servers, databases, or other componentry. The communications network 135 may also include

and be connected to a mesh network, a local network, a cloud-computing network, an IMS network, a VoIP network, a security network, a VoLTE network, a wireless network, an Ethernet network, a satellite network, a broadband network, a cellular network, a private network, a cable network, the Internet, an internet protocol network, MPLS network, a content distribution network, any network, or any combination thereof. Illustratively, servers 140, 145, and 150 are shown as being included within communications network 135. In certain embodiments, the communications network 135 may be part of a single autonomous system that is in a particular geographic region or be part of multiple autonomous systems that span several geographic regions.

[87] In certain embodiments, the functionality of the system 100 may be supported and executed by using any combination of the servers 140, 145, 150, and 160. The servers 140, 145, and 150 may reside in communications network 135, however, in certain embodiments, the servers 140, 145, 150 may reside outside communications network 135. The servers 140, 145, and 150 may provide and serve as a server service that performs the various operations and functions provided by the system 100. In certain embodiments, the server 140 may include a memory 141 that includes instructions, and a processor 142 that executes the instructions from the memory 141 to perform various operations that are performed by the server 140. The processor 142 may be hardware, software, or a combination thereof. Similarly, the server 145 may include a memory 146 that includes instructions, and a processor 147 that executes the instructions from the memory 146 to perform the various operations that are performed by the server 145. Furthermore, the server 150 may include a memory 151 that includes instructions, and a processor 152 that executes the instructions from the memory 151 to perform the various operations that are performed by the server 150. In certain embodiments, the servers 140, 145, 150, and 160 may be network servers, routers, gateways, switches, media distribution hubs, signal transfer points, service control points, service switching points, firewalls, routers, edge devices, nodes, computers, mobile devices, or any other suitable computing device, or any combination thereof. In certain embodiments, the servers 140, 145, 150 may be communicatively linked to the communications network 135, any network, any device in the system 100, or any combination thereof.

[88] The database 155 of the system 100 may be utilized to store and relay information that traverses the system 100, cache content that traverses the system 100, store data about each of the devices in the system 100 and perform any other typical functions of a database. In certain embodiments, the database 155 may be connected to or reside within the communications network 135, any other network, or a combination thereof. In certain embodiments, the database 155 may serve as a central repository for any information associated with any of the devices and information associated with the system 100. Furthermore, the database 155 may include a processor and memory or may be connected to a processor and memory to perform the various operations associated with the database 155. In certain embodiments, the database 155 may be connected to the servers 140, 145, 150, 160, the first user device 102, the second user device 111, the additional user devices, any devices in the system 100, any process of the system 100, any program of the system 100, any other device, any network, or any combination thereof.

[89] The database 155 may also store information and metadata obtained from the system 100, store metadata and other information associated with the first and second users 101, 110, store artificial intelligence models utilized in the system 100, store sensor data and/or content, store predictions made by the system 100 and/or artificial intelligence models, store confidence scores relating to predictions made, store threshold values for confidence scores, store cues generated by the system 100, store information associated with anything detected via the system 100, store information and/or content utilized to train the artificial intelligence models, store information associated with behaviors and/or actions conducted by individuals, store information associated with interactions conducted with respect to cues, store markers associated with individuals, store segments, store topics of interest, store knowledge graphs generated and/or utilized by the system 100, store user profiles associated with the first and second users 101, 110, store device profiles associated with any device in the system 100, store communications traversing the system 100, store user preferences, store information associated with any device or signal in the system 100, store information relating to patterns of usage relating to the user devices 102, 111, store any information obtained from any of the networks

in the system 100, store historical data associated with the first and second users 101, 110, store device characteristics, store information relating to any devices associated with the first and second users 101, 110, store information associated with the communications network 135, store any information generated and/or processed by the system 100, store any of the information disclosed for any of the operations and functions disclosed for the system 100 herewith, store any information traversing the system 100, or any combination thereof. Furthermore, the database 155 may be configured to process queries sent to it by any device in the system 100.

[90] In certain embodiments, the system 100 may utilize various componentry, devices, systems, and/or processes to support the operative functionality provided by the system 100. Referring now also to FIG. 2, an exemplary architectural overview of a portion of the system 100 is schematically shown. In certain embodiments, the architecture may be built on a modern microservices architecture that comprises several core components working in tandem to generate and personalize cues at a scale for any type of population, including large populations. In certain embodiments, the architecture may be utilized with container orchestration systems that facilitate deployment, management, and scaling of software containers. The container orchestration systems may be utilized for computation and data processing, distributed data processing, and deep learning in conjunction with graphs processing unit nodes. The architecture may also utilize distributed storage for data storage (e.g., cloud-base storage, etc.), and virtual machines to execute high-performance analytical databases utilized by the system 100.

[91] The system 100 may be configured to execute and/or operate in any type of cloud-computing environment and even in on-premises clusters, that may provide a high level of flexibility in deploying the functionality provided by the system 100 in diverse customer environments, including highly secured network environments. In certain embodiments, the architecture may include various additional components to facilitate the operative functionality of the system 100. The system 100 may include a subsystem 200 (depicted as a “Tower of Truth” in FIG. 2), which may include any number of databases 155 for storing user (e.g.,

participant) data, and may serve as a source of truth for all markers determined by the system 100 based on the data. The subsystem 200 may provide a common data model to ingest various types of source data, including, but not limited to, demographic data, health behavior data (e.g., physical activity, nutrition, sleep, vitals, etc.), program participation data (e.g., information relating to an individual's participation in a regimen, curriculum, prescription program, etc.), and health indicators (e.g., health screening and/or biometric data). The data models provided by the subsystem 200 may be configured to be extensible, thereby allowing new types of data to be added as such types of data become available, or as future needs arise. In certain embodiments, different customer data sources may be mapped to the subsystem 200 standard schema and ingested via extract-transform-load (ETL) processes that are implemented at each customer deployment where the system 100 and/or its operative functionality are deployed.

[92] The system 100 may include a subsystem 210 (i.e., the "Swarm" in FIG. 2) that may include machine learning and artificial intelligence capabilities. In certain embodiments, the subsystem 210 may include a collection of artificial intelligence bots that may be configured to compute, predict, and update the markers for every participant (i.e., user or individual) of the system 100 as new data arrives at the system 100, such as via subsystem 200. The subsystem 210 of microbots may keep the participants' markers up to date so that the system 100 generates the most relevant cues for the participants. The system 100 may include a subsystem 220 (i.e., the "Honeycomb" in FIG. 2) that may be configured to store the markers, which may be utilized to provide insights into each participant's current and past set of markers. In certain embodiments, the subsystem 220 may also be configured to record the pathways that each participant has taken, as their markers change over time. The pathways may, for example, include the specific actions that a participant has taken towards achieving a particular goal, which may include an identification of which cues a participant interacted with, how the participant interacted with the cues, the time of interaction with the cues, how the actions change as the markers change, along with any other information, or a combination thereof. In certain embodiments, actions may be specific things that a participant may do or perform. For example, an action may include, but is not limited to, a behavior, an activity, a process, a motion, an act, or a combination thereof.

[93] The system 100 may architecturally include a subsystem 230 (i.e., the "CueRank" in FIG. 2) that may serve as a source of digital cognitive power for the system 100. The subsystem 230 may include a deep learning-based recommender system that powers the personalized cues and pathways for each participant. Built on cutting-edge technologies, such as, but not limited to, Graph Neural Networks (GNN) and sequence-based recommender systems, the subsystem 230 recommends the next best action for each participant and selects the best corresponding cue, to a high degree of personalization. In certain embodiments, the system 100 may integrate and/or communicatively link with any number of applications, which may be third party applications, such as via any number of application programming interfaces (APIs) of the system 100. For example, the system 100 may integrate with applications that allow professionals to manage the engagement of their patient or citizen cohorts. A population health application of the system 100 may allow users to author or create cues to engage with participants and create segments to target cues to specific sub-populations of the participant pool. A clinician application may enable healthcare professionals to monitor the daily progress of the patient cohorts under their care, such as, but not limited to, monitoring daily health behaviors, and progress towards goals. As another example, the system 100 may be configured to integrate with various content delivery channels including mobile applications (e.g., via push notifications, etc.), email, simple messaging services (SMS), and other channels to deliver personalized cues to the participants of the system 100.

[94] The system 100 may also be configured to architecturally include telemetry and DevOps functionality. The system 100 may be configured to obtain telemetry data from the various applications connected to the system 100 (e.g., third party applications and/or other applications) and may learn from telemetry data that it collects as it is deployed in more populations and healthcare systems. The anonymous and privacy-controlled telemetry data may provide critical signals for engineering and data science teams to measure and understand the performance of all subsystems of the system 100, including data pipelines, machine learning models and the recommendation subsystem of the system 100. The data may be fed back into the product development cycle and may allow engineers, data scientists, and/or even software

to iterate on future versions of the system 100.

[95] Referring now also to schematic diagram 300 of FIG. 3, the system 100 may utilize the telemetry data for various purposes. For example, the telemetry data may be utilized to provide model training outputs (enables iterative tuning and improvement of the models) including improvements relating to model parameters, model performance (including loss, accuracy, precision, recall, F1 Score, AUC, etc.). Additionally, the telemetry data may be utilized with respect to attributes of knowledge graphs utilized to support the functionality of the system 100. For example, attributes of the knowledge graph (monitoring of the changes in the knowledge graph may be enabled and the models may be retrained if a significant change is observed) may be enhanced and/or improved upon using the telemetry data. Such attributes may include, but are not limited to, the number of nodes for per entity type (e.g., participants, cues), the number of relations, the total number of triples, and the number of participant-related edges and their distribution per participant (e.g., opened_cue, has_marker, in_segment). The telemetry data may include information associated with the click rate of the cues that were sent (enables monitoring of the online performance of the model). The telemetry data may also be utilized to provide information associated with experience data (completely anonymized, enables collection of data for reinforcement learning for the system 100) of the following components: state at time t, action taken at time t, reward obtained after taking the action at time t, next state at time t+1, and next action taken at time t+1.

[96] Referring now also to schematic diagram 400 of FIG. 4, an exemplary customer environment and system development environment are shown. In certain embodiments, all development may occur in the development environment. At each release cycle of the software supporting the system 100, the pre-trained models and software images may be shipped to customers using the continuous delivery pipelines. Any further model training and scoring on customer data inside the customer's environment may be performed by fully automated pipelines, without requiring a data scientist to work within the customer environment. In certain embodiments, telemetry data may be produced by the production pipelines, which may provide useful information to iterate on the next versions of the models. Referring now also to

schematic diagram 500 of FIG. 5, the same iterative development process may occur across multiple customers. Software images and models may be shipped and executed within the customer's environment. Telemetry data may be returned, and this data may be used across a spectrum of research and development activities including product planning, software engineering, data science and artificial intelligence, along with performance and bug fixes. The system 100 improves with the accumulated learnings across all system 100 customers.

[97] Referring now also to schematic diagram 600 of FIG. 6, an exemplary process flow from beginning to end of the operation of the system 100 is shown. Diagram 600 also includes information relating to variables and data utilized by the system 100 to facilitate the generation of cues to facilitate compliance and behavioral activity by participants in the system 100. The process flow may include obtaining data associated with any number of participants of the system 100. The data may include, but is not limited to, health data, behavioral data, demographic data, psychographic data, vital sign data, medical data, any type of data, or a combination thereof. The process may include determining markers from the data. The markers may include information that may identify a particular participant and/or identify characteristics associated with the participant. The process may include setting various goals and/or care plans (e.g., intrinsic goals) for each participant or obtaining the goals and/or care plans from prescriptions, regimens, instructions, and/or other plans that may have already been established for each participant. In certain embodiments, goals may be a particular aim, objective, and/or desired result sought by a participant, an advisor of the participant, or a combination thereof. The intrinsic goals may be naturally encoded destinations that may not change from customer to customer or from condition to condition. For example, goals may be to lose weight, exercise more, eat nutritious food, quit smoking, or get an annual examination for their health. Care plans may be specific regimens that may require clinical determinations, such as from a health professional. Care plans, for example, may include plans for controlling blood glucose levels, taking insulin, or testing blood glucose levels. In certain embodiments, regimens may include specific and/or prescribed courses of medical treatment, psychological treatment, any type of treatment, or a combination thereof. In certain embodiments, regimens may be specific sequences of steps that a participant may be requested to take to accomplish a

goal or objective. For example, a regimen may identify the specific types of foods that a participant should consume each day and at what time and in what order.

[98] The markers may then be utilized by the system to create and/or identify segments. Segments may comprise information that may be utilized to identify a specific sub-population of participants within a larger set of participants. For example, if a marker for a participant is that the participant is a male and another marker for a participant is that the participant has an unhealthy body mass index, an exemplary segment may be utilized to identify the participants that are male that also have unhealthy body mass indices. Based on the markers and/or segments, the system 100 may determine cues that are personalized to each individual participant in the population of participants. The cues may be tailored based to the individual and may be ranked based on the participant's likelihood to interact with the cue and/or perform a behavioral action based on the cue. The system 100 may also generate behavior graphs that may be utilized to track the specific paths/pathways that each participant takes while trying to advance towards a goal, regimen, and/or other plan. In certain embodiments, advancing towards a goal, regimen, and/or other plan may encompass taking action and/or steps towards the goal, regimen and/or other plan. In certain embodiments, the behavior graphs including the information relating to the paths/pathways may also be utilized to generate the specific cues that are selected for each participant. The system 100 may then proceed to analyze the interactions with the cues and determine what actions a participant has taken based on her interaction with the cue and if she has advanced towards the goal.

[99] Referring now also to schematic diagram 700 in FIG. 7, the diagram 700 illustrates the ranking of cues and batch-based transmission of the cues to applications for ultimate delivery to users for interaction. In Figure 7, for example, the ranking and recommendation of the top *k* best cues per user per day may be based on a batch processing job that runs daily or at any other desired time interval. Every morning (or other desired time) based on the latest end-of-day data from the day before (or other reference time), the system 100 may utilize a knowledge graph to rank all applicable cues for each participant (i.e., user) and select the top *k* cues from the ranked list. The cues may then be pushed to participants via various

communications channels, such as mobile push notifications, emails, or text messages (e.g., SMS). The cues may be delivered to participants at different times of the day (e.g., some in the morning, and others in the afternoon and evening), but, in certain embodiments, the decision of which cues to send and at what time may be made only once, such as at the start of the day when the ranking portion of the system 100 runs.

[100] Referring now also to schematic diagram 800 in FIG. 8, the diagram 800 illustrates the ability for the system 100 to generate recommendations and/or cues on-demand. In diagram 800, for example, a goal for the system 100 may be to be able to generate recommendations on-demand, and in the context of the participant's software application experience. As participants (e.g., users) are navigating the different user interface screens displayed in the application (e.g., the application may execute on first user device 102, for example), an embedded code (much like an advertising tag) may make a request to the system's 100 API for a contextual, personalized cue. In that request, the application may specify the type of cue and the context (e.g., the application may ask for a steps cue on one screen, but a nutrition cue on another). These API calls may happen as the participant navigates through the application, and the cues may be generated and recommended on the fly and returned in a matter of milliseconds. As a result, this enables the cues to be even more personalized to the participant's context (e.g., the participant's current condition, what the participant is doing, where the participant is in his pathway towards a goal, etc.). In certain embodiments, there is no limit of k cues per day – the system 100 may generate as many cues as needed, while the participant is using the application or otherwise.

[101] Referring now also to FIG. 9, an exemplary architecture 900 for facilitating the recommendation system functionality of the system 100 is shown. A generic recommendation system or a recommendation engine may utilize artificial intelligence and machine learning to help match users to items. Given a user and his context (which could be a myriad of past actions), the system 100 aims to sift through thousands and sometimes millions of items to first generate a list of candidate items of interest. The system 100 may then rank these candidate items to generate a prioritized list of recommendations. With the understanding that an output

of the system 100 is a ranked list of cues for motivating participants to perform actions, the architecture 900 illustrates how a graph neural network (GNN) can generate the rankings. Using a graph representation enables the capturing of interactions between different items (e.g., cues) and various other entities in the system 100, such as segments and markers. A neural network may then be used to learn the patterns within these interactions and how a connected participant (along with their current markers and past actions) would best match up to which cues. For example, the architecture 900 illustrates that participants (e.g., users) and their associated context may be provided to a GNN engine of the system 100 for analysis and processing. Candidate cues may be generated by the system 100 and the generated cues may be ranked for each participant according to their markers, segments, etc. In certain embodiments, information associated with generating the cues and/or the cues themselves may be obtained from an items corpus of the system 100. Referring now also to FIG. 10, an illustrative example 1000 depicting sequential information that may be utilized by GNNs of the system 100 to rank cues and/or personalize pathways is shown. For example, sequential information may be captured to rank cues and personalize pathways depending on the latest set of actions and risk markers of each participant. FIG. 10, for example, illustrates various sequential information that may be captured to rank cues and personalize pathways. For example, the system 100 may initially recommend that a participant take 1000 steps per day, however, as activities and/or behaviors are performed in response to cues, the system 100 may proceed to recommend a yoga class, then 3000 steps per day, then that the participant join a gym, and then recommend that the participant quit smoking.

[102] Referring now also to diagram 1100 of FIG. 11, exemplary pointwise loss and pairwise loss algorithms for use with recommendation systems are shown. Certain recommendation systems approach the recommendation task in two ways. In the first approach, the task is treated as a classification problem, which aims to minimize the error of predicting which items a user would like the most. Such systems may be used for predicting the rating users might give to purchased items, or the likelihood of a user clicking on a given advertisement. In the second approach, the task is converted into a ranking problem, which aims to generate a ranked list of preferences for all items that are available to the user. Since a

goal of the system 100 functionality is to select the best personalized set of cues for each user, the second approach (which ranks the cues according to user preference) is preferable as an output mechanism.

[103] Referring now also to diagram 1200 of FIG. 12, an exemplary candidate item generation process for use with the system 100 is illustrated. The steps may include taking each participant of the system and their context (e.g., markers, segments) and generating a user-item pair or if side-information is available then a trio. Each pair or trio may be transformed into a separate data instance. The system 100 may represent each data instance using features and may perform prediction and/or rankings based on user-item-knowledge interaction modeling using matrices and/or graphs. Referring now also to diagram 1300 of FIG. 13, diagram 1300 illustrates data sparsity issues that may result if the pairs or trios are treated as an independent separate data instance and each data instance is represented using features. Diagram 1400 of FIG. 14 illustrates suboptimal representations leading to bad interaction models for unseen user-item interactions (i.e., interactions with cues by participants). Diagram 1500 of FIG. 15 illustrates black-box issues where reasoning or explanations may not be doable. Diagram 1600 of FIG. 16 illustrates ways in which the foregoing issues may be addressed via the system 100.

[104] Referring now also to diagram 1700 of FIG. 17, diagram 1700 illustrates an implementation of a recommendation system for use with the system 100 where the recommendation system is implemented as a graph with deep learning. Diagrams 1800, 1900, and 2000 of FIGs. 18, 19, and 20 illustrate further details relating to implementing a recommendation system as a graph with deep learning. Diagram 2100 of FIG. 21 illustrates session-based recommendations and information relating thereto. Diagram 2200 of FIG. 22 illustrates issues relating session-based recommendations. For example, in previously existing technologies, there is difficulty in estimating user representations without adequate user behavior in one session. Another limitation of previously existing technologies is only modeling single-way transitions between consecutive items and neglecting the transitions among contexts. Also, previously existing technologies assumes that each user only accesses

one item at each timestep. Diagram 2300 of FIG. 23 illustrates possible solutions by facilitating session-based recommendation using graph neural networks. Diagram 2400 of FIG. 24 illustrates various aspects of constructing session graphs for use with the system 100. Diagram 2500 of FIG. 25 illustrates learning item embeddings on session graphs with the system 100. Diagram 2600 of FIG. 26 illustrates generating session embeddings using the system 100. Diagram 2700 of FIG. 27 illustrates various aspects of making recommendations and conducting model training based on session graphs in the system 100.

[105] Referring now also to FIG. 28, an exemplary architecture 2800 to facilitate the operative functionality of the system 100 is shown. In the architecture 2800, the architecture illustrates the determination of information associated with participants, identifying action markers and segments associated with various sub-populations of the participants, and utilizing a knowledge graph including all markers and conditions in conjunction with the recommender system to generate recommendations for next actions for participants to perform. The architecture 2800 also allows for the generation of cues and the use of a cue filter to tailor a set of cues for each participant to motivate each participant to perform the recommended activity. Depending on the participant's interaction with one or more of the set of cues tailored for the participant and/or whether or not the participant performed the activity associated with the cue, feedback associated with such interactions and/or performance of activities associated with the cues may be utilized in a user activity graph that may be utilized to update the knowledge graph utilized by the system 100 in determining markers, determining segments, recommending activities, and/or recommending cues on subsequent iterations of the processes supporting the functionality of the system 100.

[106] Referring now also to FIG. 29, an exemplary pipeline 2900 of the system 100 is shown. In certain embodiments, the pipeline 2900 may be made up of two modes of operation, a training mode and an inference mode. During the training mode (as shown in the top portion of FIG. 29), the system 100 may extract data from the subsystem 220 (i.e., the Honeycomb Database) to construct a knowledge graph of individual participants, their risk markers, available cues, and other entities. Together, these inputs may form the basis of internally

representing each individual and the set of candidate items along with their secondary information. The knowledge graph may then be used as an input into model training. There are numerous strategies internal to the system 100 that may help the training process to eventually output a trained graph-based neural network recommendation system model, such as a knowledge graph attention network model. The training process iteratively optimizes numerous parameters for prediction performance given user defined constraints.

[107] Referring now also to FIG. 30, further details relating to the exemplary pipeline 2900 are shown. In FIG. 30, the system 100 may be in an inference mode of operation. While in the inference mode of operation, the system 100 pipeline 2900 may execute on a nightly basis (or at any other desired time), using the trained model to generate a ranked list of recommended cues for each participant of the system 100. The system 100 may start by selecting a list of candidate cues for each participant. In certain embodiments, the candidate cues may be a subset of cues that can be sent to each participant. As indicated in the present disclosure, each cue generated in the system 100 may have an associated set of segments. A given cue may be selected to be a candidate cue if the cue's target segment matches a participant's behaviour segment. For example, if the cue's target segment is to target those participants with high body mass index that are male and this target segment matches the segment for the participant, the cue may be selected for that participant. In certain embodiments, other criteria for candidate cue selection may also be utilized by the system 100, such as excluding cues for which a participant has provided negative feedback, has not interacted with, and/or for which a desired action was not performed in response to the cue.

[108] Referring now also to FIG. 31, further details relating to the inference mode of operation for the pipeline 2900 are shown. For example, in FIG. 31, in addition to the candidate cue selection process, the system 100 may also incrementally update the knowledge graph during the inference operation to include changes in behavior and other data since the previous inference timestep. Next, the system 100 may split the participants into at least two groups, namely connected and disconnected participants, depending on whether they have any connected edges in the knowledge graph. In certain embodiments, participants may be

connected in the knowledge graph if they have at least one marker identified by the system 100 or if they have interacted with at least one cue. For the connected participants, the knowledge graph and the knowledge graph attention network model (e.g., artificial intelligence model for use with the system 100) may be used to generate a ranked list of recommended cues with the highest probability of being opened by a participant. In the event where a participant does not have any connections in the knowledge graph (e.g., a disconnected participant), cues may be randomly selected by the system 100 from the list of candidate cues to be sent to the participant. Because the knowledge graph is dynamic, it is possible for a connected participant to become disconnected if the participant does not have any static markers (e.g., Sex | IsMale) and the participants loses his remaining marker, segment, and cue connections due to the participant's status, such as inactivity.

[109] Referring now also to FIG. 32, still further details relating to the inference mode of operation for the pipeline 2900 of the system are shown. In FIG. 32, for example, the stage of the pipeline 2900 may be the stage at which recommended cues are sent to participants of the system 100 and interactions with the cues may be monitored by the system 100. After the recommended cues are sent to the participants, such as to first user device 102, the responses of each participant to each cue may be collected and stored by the system 100. For example, the responses may include information relating to whether the participant has opened the cue, interacted with the cue, how the participant interacted with the cue, how long the participant interacted with the cue, whether the participant gave positive feedback regarding the cue, whether the participant gave negative feedback regarding the cue, any other information associated with the cue and/or feedback associated with the cue, or a combination thereof. In certain embodiments, the responses may be stored in the subsystem 200 (i.e., the Honeycomb Database) and may be utilized in training the models of the system 100 and to facilitate inference relating to which cues may be best suited for which participants.

[110] In certain embodiments, the system 100 may utilize knowledge graph attention network models to facilitate the operative functionality provided by the system. Referring now also to FIG. 33, an exemplary knowledge graph attention network 3300 that may be utilized to

facilitate the generation of recommendations for cues for participants is illustrated. In certain embodiments, the knowledge graph attention network may comprise a knowledge graph-based recommendation method that leverages a collaboration knowledge graph (CKG) to perform recommendations. A CKG may fuse a user-item bipartite graph with an item knowledge graph to exploit the high-order connectivity between users and items, resulting in improved recommendations. The input to the model may be a CKG and the model may output a prediction probability y_ui on whether user *u* would like item *i*. The knowledge graph attention model may begin with an embedding layer where all entities and relations of the CKG are parameterized as vector representations. Next, in certain embodiments, in the attentive embedding propagation layers, which may be built upon graph convolution networks (GCN) and graph attention network (GAN), embeddings may be recursively propagated along high-order connectivity to update their representations. Each layer may be made up of three components, which may include: 1. information propagation, 2. knowledge-aware attention and 3. information aggregation. Referring now also to FIG. 34, in the prediction layer of the knowledge graph attention network 3300, the representations of a target user and item may be aggregated across all layers and may be utilized to predict the probability that the user would like the item. The objective function defined in equation 10, as shown in FIG. 34, which comprises of both the knowledge graph loss and the collaborative filtering (CF) loss is used to optimize the model.

[111] Referring now also to FIG. 35, a table 3500 is provided, which illustrates adaptations made to a knowledge graph attention network to customize it for operation with the system 100. For example, in terms of user connections that may be made via the system 100, a typical knowledge graph attention network may have users only connected to items, however, when the knowledge graph attention network is modified for the system 100, the users may not only be connect to items (e.g., cues), but also to other types of entities, such as markers and/or segments. With regard to cold start scenarios, a typical knowledge graph attention network may not be able to handle a cold start scenario and may only be able to make predictions on entities that were observed during training of the model. In contrast, the knowledge graph attention network modified for the system 100 may be configured to readily handle a cost start

scenario and generate predictions on unseen (or unobserved) entities without the need to retrain the artificial intelligence model. With regard to model parameters, the knowledge graph attention network modified for the system 100 may be configured to utilize a greater number of parameters than those with a typical knowledge graph attention network. Furthermore, with regard to model outputs, the typical knowledge graph attention network may return a probability on whether a user may like an item (e.g., a cue), however, the modified knowledge graph attention network may return a ranked list of items (e.g., cues) for a user in order of user preferences.

[112] Referring now also to FIG. 36, a schematic diagram 3600 of the modified knowledge graph attention network for use with the system 100 is shown. In certain embodiments, the main input to the modified knowledge graph attention network may be a knowledge graph, consisting, for example, of six entities and eleven relations. These entities may be a. Participant, b. Cue c. Marker, d. Segment, e. Topic, and f. Type. The 11 relationships may be – 1. opened cue, 2. in segment, 3. has marker, 4. improves to, 5, has goal, 6. has audience, 7. includes, 8. excludes, 9. in topic, 10. benefits and 11. in type. While several of the relations in the knowledge graph may remain static, the participant-related relations ('#1 opened cue', '#2 in segment' and #3 'has marker') may change daily (or at other time intervals) as participants open new cue recommendations, alter their behavior, and are identified to have or not have new set of markers, and be in or out of current set of segments. For each participant in the knowledge graph, the modified knowledge graph attention network outputs a list of the top recommended cues (e.g., three top cues) that are most likely to be opened by each participant. Referring now also to FIG. 37, a table 3700 illustrating various parameters that may be utilized with the modified knowledge graph attention network is shown. For example, there may be a total of nine model parameters that may be optimized in in the modified knowledge graph attention network, which may include: 1. entity_dim, 2. relation_dim, 3. conv_layers, 4. dropout, 5. agg_type, 6. neg_slope, 7. LR, 8. L2 and 9. train_kg. Corresponding descriptions and exemplary values have been included in the table 3700. Referring now also to FIG. 38, a table 3800 is illustrated that contains experiment results obtained from training the modified knowledge graph attention network on synthetic data that was generated. After

multiple iterations, the best model (Model 17) with a CF loss of 0.42 and train and test AUC of 0.70 and 0.71 respectively was determined.

[113] Referring now also to FIG. 39, an exemplary knowledge graph 3900 illustrating relationships between participants, cues, markers, segments, and topics is shown. For example, graph 3900 depicts exemplary entities including four participants: P_1, P_2, P_3, and P_4, three cues: C_1, C_2, and C_3: two topics: T_1 ("BMI") and T_2 ("Steps"); two segments: S_1 ("Men with Unhealty BMI") and S_2 ("Fitness Buffs"); and six markers: M_1 ("BMI | isHigh BMI"), M_2 ("Sex | IsMale"), M_3 ("BMI | IsHealthyBMI"), M_4 ("BMI | IsVeryHighBMI"); M_6 ("is10kStepper"); and M_7 ("Physical Activity | IsHighAerobicActivity). Exemplary edges of the knowledge graph 3900 may include r_0 (opened cue), r_1 (in segment), r_2 (has marker), r_3 (has audience), r_4 (has goal), r_5 (includes), and r_6 (excludes). Regardless of individual preference or health status, the system 100 is able to personalize cues for each participant. The system 100 may do so by exploiting the relationships between participants, cues, markers, segments, and topics in the knowledge graph to find the most suitable cue to recommend to each participant. Referring now also to FIG. 40, an exemplary use-case scenario using the knowledge graph 3900 is shown. For example, FIG. 40 illustrates how personalization works within the modified knowledge graph attention network using an example of two participants, participants 1 and 3. In FIG. 40, it can be observed that the participants 1 and 3 have a lot in common. In particular, participants 1 and 3 both opened cue 3 and belong to the 'Men with Unhealthy BMI' segment. Also, cue 3 is most like cue 1 because cue 3 and cue 1 share the same audience and goal of 'Men with Unhealthy BMI' and 'BMI' respectively. Since Participant 1 has opened cue 1 previously and given the similarities between cues 1 and 3 and Participants 1 and 3, the modified knowledge graph attention network of the system 100 will rank cue 1 for participant 3 very highly and recommend cue 1 to participant 3 as he may also be likely to open cue 1. The system 100 addresses various issues: 1. The system 100 ensures that higher engagement participants do not bias the system 100 in favor of cues they open, 2. cues that target more users do not get over exposure, etc.

[114] Another example of the system 100 in operation is shown in FIG. 41, which focuses

on two participants: participant 2 and participant 4. Although participant 2 has not opened any cues in the past, participant 2 is most like (i.e., similar or correlated with) participant 4 as they both fall in the 'Fitness Buffs' segment. Since participant 4 has opened cue 2 on a previous occasion, the modified knowledge graph attention network will rank cue 2 for participant 2 very highly and recommend cue 2 to participant 2 as he may also be likely to open it. A further example of the system 100 in operation is shown diagram 4200 in FIG. 42. This example features yet another strength of the system 100 in that the system 100 can adapt along with changing user behavior by recognizing the participant's connections dynamically and continue to recommend relevant cues to the participant iteratively as his connections evolve. A dynamic adaptation of the system is illustrated using an example participant and how the segments and cues she interacts with helps the system 100 determine the ranking of the cues to send to her. For example, the participant may be named Ann. Ann is a 38-year-old female with a recent HbA1C measurement of 6.2% and over the past week, Ann did a total of 25 mins of physical activity, which is low (150 mins/week may be the recommended level). As a result, Ann has two active markers 'Diabetes | IsPrediabetic' and 'Physical Activity | IsSedentary'. Because of these markers and her profile information, the system 100 determines that Ann falls into the 'Sedentary Prediabetics' segment. For simplicity, it may be assumed that Ann is unique and does not fall under any of the previously discussed system 100 scenarios. Referring now also to diagram 4300 in FIG. 43, within the system 100, there is a cue targeting the 'Sedentary Prediabetics' segment with a goal of increasing her physical activity. The ranking function of the system 100 may rank this cue very highly and recommend it to Ann because the cue stresses the importance of exercise and suggests a short 15 mins workout.

[115] Referring now also to FIG. 44, a diagram 4400 is shown that shows what occurs in the system 100 when a participant interacts with a recommended cue. As shown in FIG. 44, Ann opens the recommended cue and tries to include more physical activities in her daily routine. As a result, her HbA1C levels fell slightly to 5.9% and she clocked a total of 52 minutes of vigorous physical activity in a week. This change is also reflected in the knowledge graph where Ann loses the 'Physical Activity | IsSedentary' and 'Sedentary Prediabetics' connections and gains the 'Physical Activity | IsLowAerobicActivity' marker and 'Active

Prediabetic Adults' segment. Referring now also to FIG. 45, a diagram 4500 is shown that illustrates the recommendation of another cue based on the participant's progress. For example, the system 100 determines that Ann has progress towards her goal. As a result, the system 100 identifies and sends a cue to Ann, which was ranked highly to compliment her for lowering her HbA1C levels and remind her to keep up with her workouts. Referring now also to diagram 4600 of FIG 46, diagram 4600 illustrates the participant's further progress towards her goal based on her interaction with the most recently recommended cue. For example, Ann opens the most recently recommended cue and is motivated to exercise even more after seeing improvements in her HbA1C levels. This further lowered her HbA1C to 5.6%, which falls within the normal range, and she did a total of 95 minutes of vigorous physical activity in the past week, which caused her to earn the 'Physical Activity | IsMediumAerobicActivity' marker. Ann is now associated by the system 100 with the 'Fitness Fanatics' segment, which consists of participants with medium to high aerobic activity.

[116] Referring now also to diagram 4700 of FIG. 47, diagram 4700 illustrates how the system 100 can reduce the potential regression of a participant with respect to her goal. For example, to prevent Ann from regressing, and, instead, push her to further increase her physical activity, the system 100 may send out another cue, which was ranked highly to remind Ann of her progress and encourage her to book a workout session so that she can continue to make progress towards her goal. Referring now also to diagram 4800 of FIG. 48, Ann may have seen the benefits of keeping fit and may decide to continue towards her goal, and, as a result, Ann may open the cue from diagram 4700. After interacting with and opening the cue, Ann may schedule a high intensity interval training class. She may have enjoyed the training session and may have booked a few more classes. By participating in the extra sessions, this may have increased Ann's amount of vigorous physical activity to 160 minutes a week and, as a result, the system 100 may determine that she has gained a marker for 'Physical Activity | IsHighAerobicActivity' (i.e., M7 in diagram 4800).

[117] Based on the foregoing, the present disclosure describes how the system may recommend different cues in response to Ann's (or another participant) behavior. Referring

now also to diagram 4900 of FIG. 49, diagram 4900 illustrates that the system 100 may not only have focused on diabetes and physical activity for Ann, but also the system may personalize other aspects of Ann's health as well. For example, the system 100 may personalize Ann's journey towards health wellness by providing cues for achieving a healthier diet, cues for adjusting her sleep schedule or patterns, cues for adding more steps to her daily routine, and the like. As a result, the system 100 may be configured to optimize across multiple aspects for a particular type of goal or even set of goals.

[118] In certain embodiments, the system 100 may also be configured to generate on-demand cues for participants. Referring now also to diagram 5000 of FIG. 50, an exemplary process for generating cues on-demand is shown. For example, a potential use case scenario may occur whenever a user opens an application that is utilized to connect, interact, and/or interface with the system 100. Since the user's attention is already on the application, it may be an opportune time for the user to receive a cue from the system 100 and to read, view, or otherwise interact with the cue. As a result, in certain embodiments, opening the application may trigger a synchronization with the user's fitness tracker to fetch the user's latest health and activity metrics. Marker computation may then happen on-the-fly, the updated markers may be passed on to the modified knowledge graph attention network model via an API call and a recommended cue may be returned. Lastly, the recommended cue may be sent to the user and it may show up on the user's application interface as a notification.

[119] Referring now also to framework diagram 5100 of FIG. 51, the diagram 5100 illustrates various aspects and variables that are factored into and/or selected by the system 100 to support the operative functionality of the system 100. Referring now also to framework diagram 5200 of FIG. 52, the diagram 5200 illustrates exemplary design choices for the system 100 for the various aspects and variables that are factored into and/or selected by the system 100 to support the functionality of the system 100. In certain embodiments, for pre-processing, the system 100 may compute all relevant markers and/or segments based on a trigger event. For example, the trigger may be a time trigger (batch) and/or API call (on demand). In certain embodiments, the recommendation may be done on a batch mode or on-demand mode. In

certain embodiments, the system 100 may be configured to incremental computations where the system 100 only updates what is needed, based on dependency graph including tower data, markers, segments, and/or cues.

[120] In certain embodiments, the system 100 may determine the correct order of cues that are most effective for a participant. In order to do so, the system 100 may determine the correct order of cues that are most effective based on learning, review by SME, use as seed / pre-train model, and fine tune / retrain / adapt in new environment / population. Behavior graph may be created by SME, optimization trade-offs e.g., cost vs time, direct path to goal with low propensity vs indirect path with higher propensity. For clinical graphs, they may be less fuzzy than behavior graphs, there may be more rules and dependencies, and there may be measurements, weightages, thresholds, care plans, and goals. In certain embodiments, care plans may be the same for all people who have the same condition and stage, e.g., pre-diabetes care plan, diabetes care plan e.g., take prescription daily, do daily blood test. Individual details may be different, e.g., patient A may be prescribed daily glucose test but not patient B. In certain embodiments, the system 100 may set up reward function / goals and may combine multiple goals into a reward function. With regard to the cold start problem: a new customer may not have data for model training. In certain embodiments, the system 100 may be configured to distinguish between notifications and cues. For example, too many notifications may annoy a participant, but the system 100 can generate more cues in the application that are not necessarily used as notifications. In certain embodiments, the system 100 may also do transfer learning where there is a fixed set of markers (and other nodes) in the system knowledge graph and there may be a map from customer-configured markers to this fixed set of markers. In certain embodiments, other functionality may be provided by the system 100. For example, a digital dashboard may be provided by the system 100 to see who is not being reached by cues and/or participants who are missed from defined segments.

[121] Referring now also to diagram 5300 of FIG. 53, the diagram 5300 depicts an ability of the system 100 to utilize learned embeddings for the same graph nodes from multiple graph deployments and combining the learnings to generate new versions of knowledge graphs

utilized by the system 100. Starting with GNNs, typical outputs of GNNs may include (1) weight matrices and (2) learned embeddings of the nodes in the graph. In certain embodiments, some parts of the graph may be fixed. For example, markers, topics (goals), cue signatures, segment signatures, and/or participant signatures may be fixed. As an illustrative example, the system 100 may provide a set of markers out-of-the-box to all customers, and the markers can be baked into the core knowledge graph utilized by the system 100. As the system 100 functionality is deployed in different places and fine-tuned on customer data, the learned embeddings for these graph nodes may differ from one deployment to another. The system 100 may collect the different learned embeddings for the same graph nodes from multiple deployments and combine the learnings to improve the next version of the knowledge graph utilized by the system 100.

[122] Referring now also to diagram 5400 of FIG. 54, diagram 5400 illustrates certain portions of learned embeddings of a knowledge graph that may be transferable. In certain embodiments, at the most basic level, the taxonomy of out-of-the-box markers may be fixed and may appear in every deployment of system 100. The sub-graph circled in FIG. 54 may be common across deployments, so the learned embeddings may be transferable. In contrast, the other parts of the knowledge graph may be based on customer data and may not be easily transferrable. For example, an exemplary participant John Smith having ID 1234567 is unlikely to appear in the data of multiple customers, so the learned embedding for John Smith may not be transferrable to a different deployment. Similarly, the specific cues and segments created by customers may be unique to their specific usage context.

[123] Referring now also to diagram 5500 of FIG. 55, diagram 5500 illustrates the ability of the system 100 to create signatures for participants, cues, and segments. For example, instead of identifying individual participants, cues and segments, the system 100 can create "signatures" for them, based on their attributes. As an illustrative example, a participant can be defined as "male, 30 years old, prediabetic, smoker" instead of just as "John Smith" for the participant's signature. In that way, the system 100 may learn the embeddings of all other 30-year-old male prediabetic smokers and transfer the learnings of this type of person to other

deployments and contexts. Similarly, cues and segments may be described by a fixed taxonomy of attributes. Instead of focusing on the actual content of the cue, the system 100 may be configured to learn embeddings of cues with common attributes, such as the goal, the number of characters, and the tone, and transfer the learnings and/or insights to other cues with different content, but with similar characteristics and/or correlations. In certain embodiments, signatures may also be generated for cues, segments, and/or other variables utilized by the knowledge graph and/or system 100.

[124] Referring now also to FIG. 56, a diagram 5600 illustrating examples of determining the number of possible participant signatures based on markers in the signature are shown. One possible challenge in using signatures is ensuring k-anonymity. The more information the signature contains, the higher the probability that only one or two individuals have a signature. For example, if a participant's signature is defined by just 35 markers, there may be 181,400 unique combinations. If the population has only 500,000 individuals, then it is very likely that many or most of them have unique signatures not shared with anybody else in the population. The unique set of markers may be used to re-identify individual participants even if the data were anonymized. Therefore, this system approach of using signatures may be tuned in the context of the expected population sizes of system 100 deployments.

[125] Referring now also to FIG. 57, a diagram 5700 illustrating reinforcement learning capabilities of the system 100 is shown. For example, an approach for transfer learning borrows from the concept of experience replay in the field of Reinforcement Learning. Many reinforcement algorithms (e.g., such as the algorithms trained to play chess, the game go, or massive online multiplayer games) benefit from experiencing millions of different states of the world and figuring out how to make optimal choices leading to the highest rewards. The system 100 may leverage similar ideas and the use of reinforcement learning for recommender systems of the system 100 may be incorporated to facilitate the functionality of the system 100. In the context of the system 100, reinforcement learning may operate in the following manner: (a) Read the state of the participant (e.g., health status, recent behavior, medical history, etc.); (b) Perform an action (e.g., make a recommendation); (c) Measure the reward of taking that

action (e.g., behavior change); (d) Update the internal weights (this may be where learning takes place); (e) Repeat the process. Given enough time and data, the system 100 may learn what are the best actions (e.g., recommendations) to take for any given input state of the participant.

[126] Referring now also to diagram 5800 of FIG. 58, diagram 5800 illustrates experience data for use with reinforcement learning. In certain embodiments, reinforcement learning may be trained on "experience data" that is represented by a tuple of 5 values: S1: Initial state at time t=1; A1: Optimal action taken at state S1; R1: Observed reward after taking the action, at time t=1; S2: New state of the participant, at time t=2; and A2: Optimal action taken at state S2. In certain embodiments, each time the system 100 makes a recommendation and observes the result, the system 100 accumulates a database of experience data. Given enough time and making recommendations for large populations, the system 100 may quickly accumulate many millions of rows of such experience data so that the system 100 may utilize the experience data use for further learning and fine-tuning. For example, if the system 100 generates recommendations for 1 million people every day, then in 1 year, the system 100 may accumulate 365 million rows of experience data. The experience data may be completely anonymous. For example, there may be no user identifiers or any other identifiers in the experience data. In certain embodiments, all that may be required is a way to encode the state, action, and reward values (e.g., the user's health markers, the type of cue sent, and the behavior change).

[127] Referring now also to diagram 5900 of FIG. 59, diagram 5900 illustrates additional information relating to the use of experience data to facilitate reinforcement learning in the system 100. A benefit of collecting experience data in the described manner is that the same data schema and collection methods may be used across all deployments, and the data from multiple customers may be aggregated in a completely anonymous way. Over time, this creates a rich and comprehensive database of recommendation experiences across countries, diseases, and cultures, which may be used to train future versions of the software supporting the functionality of the system 100. Referring now also to diagram 6000 of FIG. 60, the diagram

6000 illustrates a way in which to tune and create different types of recommendation models for use with the system 100. For example, the experience database described herein may be utilized to tune and create many different recommendation models with different parameters that can be adjusted according to the need. In the equation shown in FIG. 60, Q is the model that is being learned by the system 100. The learning rate can be adjusted to influence how quickly the model adapts and adjusts as it sees new data. The discount rate influences whether the model optimizes for short-term rewards (e.g., get the participant to walk this weekend), or longer-term ones (e.g., reduce the participant's A1C metric over a time period of 12 months).

[128] In certain embodiments, the system 100 may utilize any number and/or type of artificial intelligence models to support the functionality of the system 100. In certain embodiments, an artificial intelligence model may be a file, program, module, and/or process that may be trained by the system 100 (or other system and/or subsystem described herein) to recognize certain patterns, behaviors, and/or content. For example, the artificial intelligence model(s) may be trained to detect markers for each participant of the system 100 based on data obtained for the participant. In certain embodiments, the artificial intelligence model may be, may include, and/or may utilize a Deep Convolutional Neural Network, a one-dimensional convolutional neural network, a two-dimensional convolutional neural network, a Long Short-Term Memory network, any type of machine learning system, any type of artificial intelligence system, or a combination thereof. Additionally, in certain embodiments, the artificial intelligence model may incorporate the use of any type of artificial intelligence and/or machine learning algorithms to facilitate the operation of the artificial intelligence model(s).

[129] The system 100 may train the artificial intelligence model(s) to reason and learn from data fed into the system 100 so that the model may generate and/or facilitate the generation of predictions about new data and information that is fed into the system 100 for analysis. For example, the system 100 may train an artificial intelligence model using various types of data and/or content, such as, but not limited to, images, video content, audio content, text content, augmented reality content, virtual reality content, information relating to patterns, information relating to behaviors, information relating to activities and/or occurrences,

information relating to a participant's interaction with a cue, interaction relating to a participant's feedback relating to a cue, sensor data, any data associated with the foregoing, any type of data, or a combination thereof. In certain embodiments, the content and/or data utilized to train the artificial intelligence model may be utilized to enhance the determination of markers and segments, the generation of cues, and the determination of the effectiveness of cues. As additional data and/or content is fed into the model(s) over time, the model's ability to generate and/or determine optimal cues for participants will improve and be more finely tuned.

[130] In certain embodiments, as shown in Figure 1, the system 100 may perform any of the operative functions disclosed herein by utilizing the processing capabilities of server 160, the storage capacity of the database 155, or any other component of the system 100 to perform the operative functions disclosed herein. The server 160 may include one or more processors 162 that may be configured to process any of the various functions of the system 100. The processors 162 may be software, hardware, or a combination of hardware and software. Additionally, the server 160 may also include a memory 161, which stores instructions that the processors 162 may execute to perform various operations of the system 100. For example, the server 160 may assist in processing loads handled by the various devices in the system 100, such as, but not limited to, obtaining data associated with individuals from a plurality of data sources; determining markers associated with the individuals based on the data; determining recommended actions for the individuals to perform and cues to nudge and/or motivate the individuals to perform the actions; monitoring the individuals' interactions with the cues; obtaining feedback associated with the monitoring; updating artificial intelligence models over time as feedback is obtained to facilitate determination of actions; and performing any other operations conducted by the system 100 or otherwise. In one embodiment, multiple servers 160 may be utilized to process the functions of the system 100. The server 160 and other devices in the system 100, may utilize the database 155 for storing data about the devices in the system 100 or any other information that is associated with the system 100. In one embodiment, multiple databases 155 may be utilized to store data in the system 100.

[131] Although Figures 1-62 illustrates specific example configurations of the various

components of the system 100, the system 100 may include any configuration of the components, which may include using a greater or lesser number of the components. For example, the system 100 is illustratively shown as including a first user device 102, a second user device 111, a communications network 135, a server 140, a server 145, a server 150, a server 160, and a database 155. However, the system 100 may include multiple first user devices 102, multiple second user devices 111, multiple communications networks 135, multiple servers 140, multiple servers 145, multiple servers 150, multiple servers 160, multiple databases 155, or any number of any of the other components inside or outside the system 100. Furthermore, in certain embodiments, substantial portions of the functionality and operations of the system 100 may be performed by other networks and systems that may be connected to system 100.

[132] In certain embodiments, the system 100 may execute and/or conduct the functionality as described in the method(s) that follow. As shown in Figure 61, an exemplary method 6100 for facilitating compliance and behavioral activity through signals driven via artificial intelligence is schematically illustrated. The method 6100 and/or functionality and features supporting the method 6100 may be conducted via an application of the system 100, devices of the system 100, processes of the system 100, any component of the system 100, or a combination thereof. The method 6100 may include steps for obtaining content and/or data associated with individuals, determining markers associated with the individuals, utilizing the markers, along with segments and/or topics to determine recommended actions for the individuals to perform and cues to motivate the individuals to perform the actions, and monitoring the individuals' interactions with the cues and whether the individuals have performed the actions. In certain embodiments, any number of the steps for method 6100 may involve utilizing a knowledge graph and/or one or more artificial intelligence models as described herein to implement the steps of the method 6100. At step 6102, the method 6100 may include obtaining data associated with a plurality of participants from a plurality of data sources. In certain embodiments, a plurality of participants may comprise two or more participants, five or more participants, and/or any number of participants that may participate in the system 100. The method 6100 may then include, at step 6104, determining markers

associated with the participants based on the data from the plurality of data sources. At step 6106, the method 6100 may include determining goals, regimens, and/or plans for each participant. For example, the system 100 may determine that a participant may have a health goal of reducing his triglycerides to a medically acceptable and/or safe level.

[133] At step 6108, the method 6100 may include determining paths and/or pathways that each participant has taken towards their goals, regimens, and/or plans. For example, if a particular participant has ten steps in a plan towards achieving a healthy weight and the participant has done three of the steps either in sequence or out of order, the participant's pathway towards the plan may be determined. At step 6110, the method 6100 may include determining and/or generating segments for targeting specific participants of the plurality of participants, such as for receiving certain types of cues. At step 6112, the method 6100 may include determining actions and/or behaviors for each participant to perform to advance towards a goal, regimen, plan, or a combination thereof. At step 6114, the method 6100 may include generating and/or obtaining cues based on the markers, goals, pathways, segments, topics (e.g., topics associated with plans, goals, regiments, etc.), and/or actions. At step 6116, the method 6100 may include ranking the cues for each participant based on the likelihood of each participant interacting with each cue, based on a probability of each participant performing the action, based on a probability of each participant advancing towards the goals, regimens, and/or plans based on interaction with a cue, or a combination thereof. For example, the cue with the highest probability for interaction may be the highest ranked cue in the list. In certain embodiments, any number of cues may be provided in the ranked list.

[134] At step 6118, the method 6100 may include facilitating delivery of one or more cues from the ranked list to each participant. In certain embodiments, the delivery of the cues may entail providing the cues to an application that a participant is interacting with, such as via a computing device (e.g., first user device 102). In certain embodiments, the delivery of the cues may entail making the cues accessible via an application programming interface that is configured to facilitate communicative coupling between the system 100 and any number of third-party applications and/or systems. In certain embodiments, the delivery of the cues may

entail pushing content corresponding to the cues directly to the devices of each participant. At step 6120, the method 6100 may include obtaining and analyzing telemetry data that provides information relating to the interaction with the cues by each participant. For example, the telemetry data may indicate whether a participant interacted with a cue, how long the participant interacted with the cue, how the participant interacted with the cue, whether the participant performed the activity that the cue was intended to promote, whether the participant performed a different activity than the activity that the cue was intended to promote, whether there was positive or negative feedback associated with the cue, any other information relating to the cue and/or activity, or a combination thereof. At step 6122, the method 6100 may include utilizing the telemetry data to update artificial intelligence models and/or knowledge graphs to enhance the cue generation and ranking process on subsequent operations of the method 6100, such as the next time each participant wants to progress towards a goal. In certain embodiments, the method 6100 may further incorporate any of the features and functionality described for the system 100, any other method disclosed herein, or as otherwise described herein.

[135] In certain embodiments, the present disclosure may further include a system that includes a memory that stores instructions and a processor that executes the instructions to perform operations. In certain embodiments, an operation may include obtaining data associated with a plurality of participants from a plurality of data sources. Another operation may include determining markers associated with each participant of the plurality of participants based on the data associated with the plurality of participants. A further operation may include determining goals, regimens, plans, or a combination thereof, for each participant of the plurality of participants. In certain embodiments, the system may also perform an operation that includes determining paths that each participant has towards the goals, regimens, plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes determining an action for each participant to perform to advance each participant towards the goals, regimens, plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes generating a plurality of cues for each participant based on the markers, the goals, the regimens, the plans, the paths, the action,

segments, topics, or a combination thereof. In certain embodiments, the system may perform an operation that includes ranking the plurality of cues for each participant based on a probability of each participant interacting with the plurality of cues. Furthermore, the system may perform an operation that includes facilitating delivery of at least one cue of the plurality of cues to each participant. Still further, the system may perform an operation that includes analyzing telemetry data providing information associated with interaction with the plurality of cues by each participant. Moreover, the system may perform an operation that includes updating an artificial intelligence model utilized to facilitate enhanced generation of the plurality of cues on a subsequent occasion.

[136] In certain embodiments, the system may perform an operation that includes analyzing telemetry data providing information associated with interaction with the plurality of cues by any one or more of the participants of the plurality of participants. In certain embodiments, the system may perform an operation that includes updating, based on the telemetry data, one or more artificial intelligence models utilized to facilitate enhanced generation of a next plurality of cues for motivating a next action to take to advance towards the goals, regimens, plans, or a combination thereof, determination of markers, determination of actions to be performed, ranking the cues, and/or performing any other operations of the system. In certain embodiments, the system may perform an operation that includes determining the segments to target a set of the plurality of participants for at least one of the plurality of cues. In certain embodiments, the system may perform an operation that includes generating a behavior graph to track the paths that each of the plurality of participants take towards advancing towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes generating at least a portion of the cues, a next cue to motivate a next action for each participant to perform, or a combination thereof, based on the behavior graph.

[137] In certain embodiments, the system may perform an operation that includes determining a next action for each participant to perform to advance towards the goals, the regimens, the plans, or a combination thereof, based on one or more interactions with one or

more cues of the plurality of cues by one or more participants of the plurality of participants. In certain embodiments, the system may perform an operation that includes determining a next cue to motivate the one or more participants to perform the next action to advance towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes monitoring a progress of the one or more participants of the plurality of participants in advancing towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes obtaining additional data from one or more participants of the plurality of participants. In certain embodiments, the system may perform an operation that includes updating the one or more marks based on the additional data.

[138] In certain embodiments, the system may perform an operation that includes determining a top ranked cue of the plurality of cues to recommend to each participant of the plurality of participants. In certain embodiments, the system may perform an operation that includes determining an order in which the plurality of cues are to be delivered to each participant to motivate each participant to perform the action to advance each participant towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes receiving feedback from one or more participants of the plurality of participants. In certain embodiments, the system may perform an operation that includes modifying the plurality of cues, generating new cues, or a combination thereof, based on the feedback. In certain embodiments, the system may perform an operation that includes determining whether one or more participants of the plurality of participants has regressed in advancing towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the system may perform an operation that includes selecting a different cue from the plurality of cues to deliver to each participant that has regressed in advancing towards the goals, the regimens, the plans, or a combination thereof.

[139] In certain embodiments, a further method according to embodiments of the present disclosure is provided. In certain embodiments, the method may be performed utilizing any of the components of the system 100, such as processors and/or memories of the system 100. In

certain embodiments, the method may include obtaining data associated with a plurality of participants from a plurality of data sources. In certain embodiments, the method may include determining one or more marker associated with each participant of the plurality of participants based on the data associated with the plurality of participants. In certain embodiments, the method may include determining goals, regimens, plans, or a combination thereof, for each participant of the plurality of participants. In certain embodiments, the method may include determining paths that each participant has towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the method may include determining an action for each participant to perform to advance each participant towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the method may include generating a plurality of cues for each participant based on the one or more markers, the goals, the regimens, the plans, the paths, the action, segments, topics, or a combination thereof. In certain embodiments, the method may include facilitating delivery of one or more cues of the plurality of cues to each participant.

[140] In certain embodiments, the method may include ranking the plurality of cues for each participant based on a probability of each participant interacting with the plurality of cues, a probability of each participant performing the action, or a combination thereof. In certain embodiments, the method may include selecting one or more cues from the plurality of cues to deliver to each participant based on the ranking. In certain embodiments, the method may include updating, based on telemetry data associated with interaction with the one or more cues, an artificial intelligence model to enhance determination of one or more future markers, one or more future goals, regimen, or plan, one or more future cues, or a combination thereof. In certain embodiments, the method may include predicting, by utilizing an artificial intelligence model, which cue of the plurality of cues each participant will interact with to advance towards the goals, the regimens, the plans, or a combination thereof. In certain embodiments, the method may include monitoring each response by each participant to the one or more cues. In certain embodiments, the method may include modifying a knowledge graph associated with an artificial intelligence model for generating the cues based on each response. In certain embodiments, the method may include determining whether the one or more cues has

been interacted with by one or more participants of the plurality of participants. In certain embodiments, the method may include selecting a select a different cue from the plurality of cues to deliver to the one or more participants if the one or more participants is determined to have not interacted with the one or more cues.

[141] In certain embodiments, a non-transitory computer readable medium comprising instructions, which, when loaded and executed by a processor may configure the processor to perform operations. In certain embodiments, the operations may include obtaining data associated with a plurality of participants from a plurality of data sources; determining one or more markers associated with each participant of the plurality of participants based on the data associated with the plurality of participants; determine goals, regimens, plans, or a combination thereof, for each participant of the plurality of participants; determining paths that each participant has towards the goals, the regimens, the plans, or a combination thereof; determining an action for each participant to perform to advance each participant towards the goals, the regimens, the plans, or a combination thereof; generate a plurality of cues for each participant based on the at least one marker, the goals, the regimens, the plans, the paths, the action, segments, topics, or a combination thereof; and facilitate delivery of at least one cue of the plurality of cues to each participant

[142] The systems and methods disclosed herein may include still further functionality and features. For example, the operative functions of the system 100 and method may be configured to execute on a special-purpose processor specifically configured to carry out the operations provided by the system 100 and method. In certain embodiments, the operative features and functionality provided by the system 100 and method may be utilized to improve the efficiency of computing devices that are being utilized to facilitate the functionality provided by the system 100 and the various methods disclosed herein. For example, by training the system 100 over time based on data and/or other information provided and/or generated in the system 100, a reduced amount of computer operations may be performed by the devices in the system 100 using the processors and memories of the system 100 than compared to traditional methodologies. In such a context, less processing power needs to be

utilized because the processors and memories do not need to be dedicated for processing. As a result, there are substantial savings in the usage of computer resources by utilizing the software, techniques, and algorithms provided in the present disclosure. In certain embodiments, various operative functionality of the system 100 may be configured to execute on one or more graphics processors and/or application specific integrated processors.

[143] In certain embodiments, various functions and features of the system 100 and methods may operate without any human intervention and may be conducted entirely by computing devices. In certain embodiments, for example, numerous computing devices may interact with devices of the system 100 to provide the functionality supported by the system 100. Additionally, in certain embodiments, the computing devices of the system 100 may operate continuously and without human intervention to reduce the possibility of errors being introduced into the system 100. In certain embodiments, the system 100 and methods may also provide effective computing resource management by utilizing the features and functions described in the present disclosure. For example, in certain embodiments, devices in the system 100 may transmit signals indicating that only a specific quantity of computer processor resources (e.g. processor clock cycles, processor speed, etc.) may be devoted to training the artificial intelligence model(s), generating recommended actions and/or cues, generating predictions relating to cues and/or actions, and/or performing any other operation conducted by the system 100, or any combination thereof. For example, the signal may indicate a number of processor cycles of a processor that may be utilized to update and/or train an artificial intelligence model, and/or specify a selected amount of processing power that may be dedicated to generating or facilitating any of the operations performed by the system 100. In certain embodiments, a signal indicating the specific amount of computer processor resources or computer memory resources to be utilized for performing an operation of the system 100 may be transmitted from the first and/or second user devices 102, 111 to the various components of the system 100.

[144] In certain embodiments, any device in the system 100 may transmit a signal to a memory device to cause the memory device to only dedicate a selected amount of memory

resources to the various operations of the system 100. In certain embodiments, the system 100 and methods may also include transmitting signals to processors and memories to only perform the operative functions of the system 100 and methods at time periods when usage of processing resources and/or memory resources in the system 100 is at a selected value. In certain embodiments, the system 100 and methods may include transmitting signals to the memory devices utilized in the system 100, which indicate which specific sections of the memory should be utilized to store any of the data utilized or generated by the system 100. In certain embodiments, the signals transmitted to the processors and memories may be utilized to optimize the usage of computing resources while executing the operations conducted by the system 100. As a result, such functionality provides substantial operational efficiencies and improvements over existing technologies.

[145] Referring now also to Figure 62, at least a portion of the methodologies and techniques described with respect to the exemplary embodiments of the system 100 can incorporate a machine, such as, but not limited to, computer system 6200, or other computing device within which a set of instructions, when executed, may cause the machine to perform any one or more of the methodologies or functions discussed above. The machine may be configured to facilitate various operations conducted by the system 100. For example, the machine may be configured to, but is not limited to, assist the system 100 by providing processing power to assist with processing loads experienced in the system 100, by providing storage capacity for storing instructions or data traversing the system 100, or by assisting with any other operations conducted by or within the system 100. As another example, the computer system 6200 may assist with generating models associated with generating predictions relating to which cues would suit which individuals within a population as individual behaviors change over time. As another example, the computer system 6200 may assist with updating the models over time based on data, intelligence, and analyses generated, obtained, and/or accessed by the system 100.

[146] In some embodiments, the machine may operate as a standalone device. In some embodiments, the machine may be connected (e.g., using communications network 135,

another network, or a combination thereof) to and may assist with operations performed by other machines and systems, such as, but not limited to, the first user device 102, the second user device 111, the server 140, the server 145, the server 150, the database 155, the server 160, any other system, program, and/or device, or any combination thereof. The machine may be connected with any component in the system 100. In a networked deployment, the machine may operate in the capacity of a server or a client user machine in a server-client user network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine may comprise a server computer, a client user computer, a personal computer (PC), a tablet PC, a laptop computer, a desktop computer, a control system, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while a single machine is illustrated, the term "machine" shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[147] The computer system 6200 may include a processor 6202 (e.g., a central processing unit (CPU), a graphics processing unit (GPU), or both), a main memory 6204 and a static memory 6206, which communicate with each other via a bus 6208. The computer system 6200 may further include a video display unit 6210, which may be, but is not limited to, a liquid crystal display (LCD), a flat panel, a solid state display, or a cathode ray tube (CRT). The computer system 6200 may include an input device 6212, such as, but not limited to, a keyboard, a cursor control device 6214, such as, but not limited to, a mouse, a disk drive unit 6216, a signal generation device 6218, such as, but not limited to, a speaker or remote control, and a network interface device 6220.

[148] The disk drive unit 6216 may include a machine-readable medium 6222 on which is stored one or more sets of instructions 6224, such as, but not limited to, software embodying any one or more of the methodologies or functions described herein, including those methods illustrated above. The instructions 6224 may also reside, completely or at least partially, within the main memory 6204, the static memory 6206, or within the processor 6202, or a

combination thereof, during execution thereof by the computer system 6200. The main memory 6204 and the processor 6202 also may constitute machine-readable media.

[149] Dedicated hardware implementations including, but not limited to, application specific integrated circuits, programmable logic arrays and other hardware devices can likewise be constructed to implement the methods described herein. Applications that may include the apparatus and systems of various embodiments broadly include a variety of electronic and computer systems. Some embodiments implement functions in two or more specific interconnected hardware modules or devices with related control and data signals communicated between and through the modules, or as portions of an application-specific integrated circuit. Thus, the example system is applicable to software, firmware, and hardware implementations.

[150] In accordance with various embodiments of the present disclosure, the methods described herein are intended for operation as software programs running on a computer processor. Furthermore, software implementations can include, but not limited to, distributed processing or component/object distributed processing, parallel processing, or virtual machine processing can also be constructed to implement the methods described herein.

[151] The present disclosure contemplates a machine-readable medium 6222 containing instructions 6224 so that a device connected to the communications network 135, another network, or a combination thereof, can send or receive voice, video or data, and communicate over the communications network 135, another network, or a combination thereof, using the instructions. The instructions 6224 may further be transmitted or received over the communications network 135, another network, or a combination thereof, via the network interface device 6220.

[152] While the machine-readable medium 6222 is shown in an example embodiment to be a single medium, the term "machine-readable medium" should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches

and servers) that store the one or more sets of instructions. The term "machine-readable medium" shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that causes the machine to perform any one or more of the methodologies of the present disclosure.

[153] The terms "machine-readable medium," "machine-readable device," or "computer-readable device" shall accordingly be taken to include, but not be limited to: memory devices, solid-state memories such as a memory card or other package that houses one or more read-only (non-volatile) memories, random access memories, or other re-writable (volatile) memories; magneto-optical or optical medium such as a disk or tape; or other self-contained information archive or set of archives is considered a distribution medium equivalent to a tangible storage medium. The "machine-readable medium," "machine-readable device," or "computer-readable device" may be non-transitory, and, in certain embodiments, may not include a wave or signal per se. Accordingly, the disclosure is considered to include any one or more of a machine-readable medium or a distribution medium, as listed herein and including art-recognized equivalents and successor media, in which the software implementations herein are stored.

[154] The illustrations of arrangements described herein are intended to provide a general understanding of the structure of various embodiments, and they are not intended to serve as a complete description of all the elements and features of apparatus and systems that might make use of the structures described herein. Other arrangements may be utilized and derived therefrom, such that structural and logical substitutions and changes may be made without departing from the scope of this disclosure. Figures are also merely representational and may not be drawn to scale. Certain proportions thereof may be exaggerated, while others may be minimized. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

[155] Thus, although specific arrangements have been illustrated and described herein, it should be appreciated that any arrangement calculated to achieve the same purpose may be

substituted for the specific arrangement shown. This disclosure is intended to cover any and all adaptations or variations of various embodiments and arrangements of the invention. Combinations of the above arrangements, and other arrangements not specifically described herein, will be apparent to those of skill in the art upon reviewing the above description. Therefore, it is intended that the disclosure not be limited to the particular arrangement(s) disclosed as the best mode contemplated for carrying out this invention, but that the invention will include all embodiments and arrangements falling within the scope of the appended claims.

[156] The foregoing is provided for purposes of illustrating, explaining, and describing embodiments of this invention. Modifications and adaptations to these embodiments will be apparent to those skilled in the art and may be made without departing from the scope or spirit of this invention. Upon reviewing the aforementioned embodiments, it would be evident to an artisan with ordinary skill in the art that said embodiments can be modified, reduced, or enhanced without departing from the scope and spirit of the present disclosure.

CLAIMS

We claim:

1. A system, comprising:

a memory that stores instructions; and

a processor that executes the instructions to configure the processor to:

obtain data associated with a plurality of participants from a plurality of data sources;

determine at least one marker associated with each participant of the plurality of participants based on the data associated with the plurality of participants;

determine goals, regimens, plans, or a combination thereof, for each participant of the plurality of participants;

determine paths that each participant has towards the goals, the regimens, the plans, or a combination thereof;

determine an action for each participant to perform to advance each participant towards the goals, the regimens, the plans, or a combination thereof;

generate a plurality of cues for each participant based on the at least one marker, the goals, the regimens, the plans, the paths, the action, segments, topics, or a combination thereof;

rank the plurality of cues for each participant based on a probability of each participant interacting with the plurality of cues, a probability of each participant performing the action, or a combination thereof; and

facilitate delivery of at least one cue of the plurality of cues to each participant.

2. The system of claim 1, wherein the processor is further configured to:

analyze telemetry data providing information associated with interaction with the plurality of cues by each participant; and

update, based on the telemetry data, an artificial intelligence model utilized to facilitate enhanced generation of a next plurality of cues for motivating a next action to take to advance towards the goals, regimens, plans, or a combination thereof.

3. The system of claim 1, wherein the processor is further configured to determine the segments to target a set of the plurality of participants for at least one of the plurality of cues.

4. The system of claim 1, wherein the processor is further configured to generate a behavior graph to track the paths that each of the plurality of participants take towards advancing towards the goals, the regimens, the plans, or a combination thereof.

5. The system of claim 4, wherein the processor is further configured to generate at least a portion of the cues, a next cue to motivate a next action for each participant to perform, or a combination thereof, based on the behavior graph.

6. The system of claim 1, wherein the processor is further configured to:

determine a next action for each participant to perform to advance towards the goals, the regimens, the plans, or a combination thereof, based on at least one interaction with at least one cue of the plurality of cues by at least one participant of the plurality of participants; and

determine a next cue to motivate the at least one participant to perform the next action to advance towards the goals, the regimens, the plans, or a combination thereof.

7. The system of claim 1, wherein the processor is further configured to monitor a progress of at least one participant of the plurality of participants in advancing towards the goals, the regimens, the plans, or a combination thereof.

8. The system of claim 1, wherein the processor is further configured to:

obtain additional data from at least one participant of the plurality of participants; and

update the at least one marker based on the additional data.

9. The system of claim 1, wherein the processor is further configured to determine a top ranked cue of the plurality of cues to recommend to each participant of the plurality of participants.

10 The system of claim 1, wherein the processor is further configured to determine an order in which the plurality of cues are to be delivered to each participant to motivate each participant to perform the action to advance each participant towards the goals, the regimens, the plans, or a combination thereof.

11. The system of claim 1, wherein the processor is further configured to:

receive feedback from at least one participant of the plurality of participants; and

modify the plurality of cues, generate new cues, or a combination thereof, based on the feedback.

12. The system of claim 1, wherein the processor is further configured to:

determine whether at least one participant of the plurality of participants has regressed in advancing towards the goals, the regimens, the plans, or a combination thereof; and

select a different cue from the plurality of cues to deliver to the at least one participant.

13. A method, comprising:

obtaining data associated with a plurality of participants from a plurality of data sources;

determining, by utilizing instructions from a memory that are executed by a processor, at least one marker associated with each participant of the plurality of participants based on the data associated with the plurality of participants;

determining, by utilizing the instructions from the memory that are executed by the processor, goals, regimens, plans, or a combination thereof, for each participant of the plurality of participants;

determining paths that each participant has towards the goals, the regimens, the plans, or a combination thereof;

determining, by utilizing the instructions from the memory that are executed by the processor, an action for each participant to perform to advance each participant towards the goals, the regimens, the plans, or a combination thereof;

generating, by utilizing the instructions from the memory that are executed by the processor, a plurality of cues for each participant based on the at least one marker, the goals, the regimens, the plans, the paths, the action, segments, topics, or a combination thereof; and

facilitating delivery of at least one cue of the plurality of cues to each participant.

14. The method of claim 13, further comprising ranking the plurality of cues for each participant based on a probability of each participant interacting with the plurality of cues, a probability of each participant performing the action, or a combination thereof.

15. The method of claim 14, further comprising selecting at least one cue from the plurality of cues to deliver to each participant based on the ranking.

16. The method of claim 14, further comprising updating, based on telemetry data associated with interaction with the at least one cue, an artificial intelligence model to enhance determination of at least one future marker, at least one future goal, regimen, or plan, at least one future cue, or a combination thereof.

17. The method of claim 14, further comprising predicting, by utilizing an artificial intelligence model, which cue of the plurality of cues each participant will interact with to advance towards the goals, the regimens, the plans, or a combination thereof.

18. The method of claim 14, further comprising:

monitoring each response by each participant to the at least one cue; and

modifying a knowledge graph associated with an artificial intelligence model for generating the cues based on each response.

19. The method of claim 14, further comprising:

determining whether the at least one cue has been interacted with by at least one participant of the plurality of participants; and

selecting a select a different cue from the plurality of cues to deliver to the at least one participant if the at least one participant is determined to have not interacted with the at least one cue.

20. A non-transitory computer readable medium comprising instructions, which, when loaded and executed by a processor, configures the processor to:

obtain data associated with a plurality of participants from a plurality of data sources;

determine at least one marker associated with each participant of the plurality of participants based on the data associated with the plurality of participants;

determine goals, regimens, plans, or a combination thereof, for each participant of the plurality of participants;

determine paths that each participant has towards the goals, the regimens, the plans, or a combination thereof;

determine an action for each participant to perform to advance each participant towards the goals, the regimens, the plans, or a combination thereof;

generate a plurality of cues for each participant based on the at least one marker, the goals, the regimens, the plans, the paths, the action, segments, topics, or a combination thereof; and

facilitate delivery of at least one cue of the plurality of cues to each participant.