



(51) International Patent Classification:

*G09B 19/00* (2006.01)      *G06F 3/01* (2006.01)

(21) International Application Number:

PCT/US2018/013429

(22) International Filing Date:

12 January 2018 (12.01.2018)

(25) Filing Language:

English

(26) Publication Language:

English

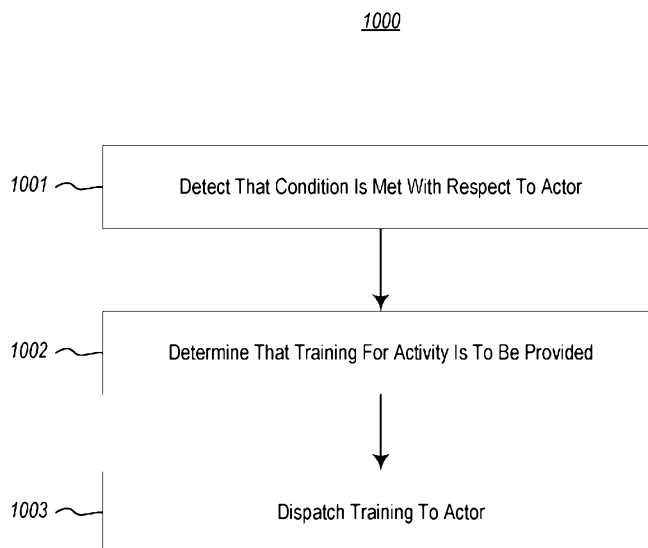
(30) Priority Data:

62/447,825      18 January 2017 (18.01.2017)      US  
15/436,683      17 February 2017 (17.02.2017)      US

(71) Applicant: **MICROSOFT TECHNOLOGY LICENSING, LLC** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **MITAL, Vijay**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **ABRAHAM, Robin**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **ZHU, Victor**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **DU, Liang**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **ZHOU, Ning**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SHARMA, Pramod Kumar**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **CHAKRABORTY, Ishani**; MICROSOFT TECHNOLOGY LICENSING,

(54) Title: AUTOMATED ACTIVITY-TIME TRAINING



**Figure 10**

(57) **Abstract:** Automatically training an actor upon the occurrence of a physical condition with respect to that actor. Upon detecting that the actor has the physical condition (e.g., is engaging in or is about to engage in a physical activity), the system determines that training is to be provided for that activity. Upon determining that training is to be provided, the system automatically dispatches training. For instance, the system might cause a human or robot to be dispatched to the actor to show the actor how to perform the activity. Alternatively or instead, a representation of a signal segment may be dispatched to the actor. The representation providing the training to the actor may include a similar target of work to what the actor is presently targeting by the activity. The representation may also include a representation of a person that engaged in the activity properly previously.

LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) **Agent:** MINHAS, Sandip S. et al.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

**Published:**

- *with international search report (Art. 21(3))*

## AUTOMATED ACTIVITY-TIME TRAINING

### BACKGROUND

[0001] Computing systems and associated networks have greatly revolutionized our world. At first, computing systems were only able to perform simple tasks. However, as processing power has increased and become increasingly available, the complexity of tasks performed by a computing system has greatly increased. Likewise, the hardware complexity and capability of computing systems has greatly increased, as exemplified with cloud computing that is supported by large data centers.

[0002] For a long period of time, computing systems just did essentially what they were told by their instructions or software. However, software and the employment of hardware is becoming so advanced that computing systems are now, more than ever before, capable of some level of decision making at higher levels. At present, in some respects, the level of decision making can approach, rival, or even exceed the capability of the human brain to make decisions. In other words, computing systems are now capable of employing some level of artificial intelligence.

[0003] One example of artificial intelligence is the recognition of external stimuli from the physical world. For instance, voice recognition technology has improved greatly allowing for high degree of accuracy in detecting words that are being spoken, and even the identity of the person that is speaking. Likewise, computer vision allows computing systems to automatically identify objects within a particular picture or frame of video, or recognize human activity across a series of video frames. As an example, face recognition technology allows computing systems to recognize faces, and activity recognition technology allows computing systems to know whether two proximate people are working together.

[0004] Each of these technologies may employ deep learning (Deep Neural Network-based and reinforcement-based learning mechanisms) and machine learning algorithms to learn from experience what is making a sound, and objects or people that are within an image, thereby improving accuracy of recognition over time. In the area of recognizing objects within a more complex imaged scene with large numbers of visual distractions, advanced computer vision technology now exceeds the capability of a human being to quickly and accurately recognize objects of interest within that scene. Hardware, such as matrix transformation hardware in conventional graphical processing units (GPUs), may

also contribute to the rapid speed in object recognition in the context of deep neural networks.

[0005] The subject matter claimed herein is not limited to embodiments that solve any disadvantages or that operate only in environments such as those described above. Rather,  
5 this background is only provided to illustrate one exemplary technology area where some embodiments described herein may be practiced.

### **BRIEF SUMMARY**

[0006] At least some embodiments described herein relate automatically training an actor upon some condition with respect to that actor. For instance, that condition might be  
10 that the actor is performing or is about to perform an activity, or that the actor is at a physical location. Upon detecting that the condition with respect to the actor is satisfied, the system determines that training is to be provided for an activity. As an example, such a determination that the actor is engaging improperly in the activity. However, such a determination may be based on any factors or training policy.

15 [0007] Upon determining that training is to be provided, the system automatically provides training to the actor. For instance, a representation of a signal segment may be dispatched to the actor, or a human being or robot may be dispatched to the actor to show the actor how to do something. If a signal representation, the representation provides the training to the actor may include a similar target of work to what the actor is presently  
20 targeting by the activity. The representation may also include a representation of a person that engaged in the activity properly previously.

[0008] This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is  
25 it intended to be used as an aid in determining the scope of the claimed subject matter.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0009] In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific  
30 embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0010] Figure 1 illustrates an example computer system in which the principles described herein may be employed;

[0011] Figure 2 illustrates an environment in which the principles described herein may operate, which includes a physical space that includes multiple physical entities and multiple sensors, a recognition component that senses features of physical entities within the physical space, and a feature store that stores sensed features of such physical entities, such that computation and querying may be performed against those features;

[0012] Figure 3 illustrates a flowchart of a method for tracking physical entities within a location and may be performed in the environment of Figure 2;

[0013] Figure 4 illustrates an entity tracking data structure that may be used to assist in performing the method of Figure 3, and which may be used to later perform queries on the tracked physical entities; Figure 5 illustrates a flowchart of a method for efficiently rendering signal segments of interest;

[0014] Figure 6 illustrates a flowchart of a method for controlling creation of or access to information sensed by one or more sensors in a physical space;

[0015] Figure 7 illustrates a recurring flow showing that in addition to creating a computer-navigable graph of sensed features in the physical space, there may also be pruning of the computer-navigable graph to thereby keep the computer-navigable graph of the real world at a manageable size;

[0016] Figure 8 illustrates a flowchart of a method for sharing at least a portion of a signal segment;

[0017] Figure 9 illustrates a flowchart of a method for automatically generating a narration of what is happening in a signal segment; and

[0018] Figure 10 illustrates a flowchart of a method for automatically training an actor that upon a particular condition occurring with respect to that actor.

### **DETAILED DESCRIPTION**

[0019] At least some embodiments described herein relate automatically training an actor upon some condition with respect to that actor. For instance, that condition might be that the actor is performing or is about to perform an activity, or that the actor is at a physical location. Upon detecting that the condition with respect to the actor is satisfied, the system determines that training is to be provided for an activity. As an example, such a determination that the actor is engaging improperly in the activity. However, such a determination may be based on any factors or training policy.

[0020] Upon determining that training is to be provided, the system automatically provides training to the actor. For instance, a representation of a signal segment may be dispatched to the actor, or a human being or robot may be dispatched to the actor to show the actor how to do something. If a signal representation, the representation provides the training to the actor may include a similar target of work to what the actor is presently targeting by the activity. The representation may also include a representation of a person that engaged in the activity properly previously.

[0021] Because the principles described herein operate in the context of a computing system, a computing system will be described with respect to Figure 1. Then, the principles of the foundation upon which ambient computing may be performed will then be described with respect to Figures 2 through 4. The obtaining of signal segments from the computer-navigable graph will then be described with respect to Figure 5. Thereafter, the application of security in the context of ambient computing will be described with respect to Figure 6. Finally, the managing of the size of the computer-navigable graph will be described with respect to Figure 7. Then three implementations that use the semantic understanding provided by the computer-navigable graph (also called herein a “physical graph” will be described with respect to Figures 8 through 10.

[0022] Computing systems are now increasingly taking a wide variety of forms. Computing systems may, for example, be handheld devices, appliances, laptop computers, desktop computers, mainframes, distributed computing systems, datacenters, or even devices that have not conventionally been considered a computing system, such as wearables (e.g., glasses, watches, bands, and so forth). In this description and in the claims, the term “computing system” is defined broadly as including any device or system (or combination thereof) that includes at least one physical and tangible processor, and a physical and tangible memory capable of having thereon computer-executable instructions that may be executed by a processor. The memory may take any form and may depend on the nature and form of the computing system. A computing system may be distributed over a network environment and may include multiple constituent computing systems.

[0023] As illustrated in Figure 1, in its most basic configuration, a computing system typically includes at least one hardware processing unit 102 and memory 104. The memory 104 may be physical system memory, which may be volatile, non-volatile, or some combination of the two. The term “memory” may also be used herein to refer to non-volatile mass storage such as physical storage media. If the computing system is distributed, the processing, memory and/or storage capability may be distributed as well.

[0024] The computing system 100 has thereon multiple structures often referred to as an “executable component”. For instance, the memory 104 of the computing system 100 is illustrated as including executable component 106. The term “executable component” is the name for a structure that is well understood to one of ordinary skill in the art in the field of computing as being a structure that can be software, hardware, or a combination thereof. For instance, when implemented in software, one of ordinary skill in the art would understand that the structure of an executable component may include software objects, routines, methods that may be executed on the computing system, whether such an executable component exists in the heap of a computing system, or whether the executable component exists on computer-readable storage media.

[0025] In such a case, one of ordinary skill in the art will recognize that the structure of the executable component exists on a computer-readable medium such that, when interpreted by one or more processors of a computing system (e.g., by a processor thread), the computing system is caused to perform a function. Such structure may be computer-readable directly by the processors (as is the case if the executable component were binary). Alternatively, the structure may be structured to be interpretable and/or compiled (whether in a single stage or in multiple stages) so as to generate such binary that is directly interpretable by the processors. Such an understanding of example structures of an executable component is well within the understanding of one of ordinary skill in the art of computing when using the term “executable component”.

[0026] The term “executable component” is also well understood by one of ordinary skill as including structures that are implemented exclusively or near-exclusively in hardware, such as within a field programmable gate array (FPGA), an application specific integrated circuit (ASIC), or any other specialized circuit. Accordingly, the term “executable component” is a term for a structure that is well understood by those of ordinary skill in the art of computing, whether implemented in software, hardware, or a combination. In this description, the term “component” may also be used. As used in this description and in the case, this term (regardless of whether the term is modified with one or more modifiers) is also intended to be synonymous with the term “executable component” or be specific types of such an “executable component”, and thus also have a structure that is well understood by those of ordinary skill in the art of computing.

[0027] In the description that follows, embodiments are described with reference to acts that are performed by one or more computing systems. If such acts are implemented in software, one or more processors (of the associated computing system that performs the

act) direct the operation of the computing system in response to having executed computer-executable instructions that constitute an executable component. For example, such computer-executable instructions may be embodied on one or more computer-readable media that form a computer program product. An example of such an operation  
5 involves the manipulation of data.

**[0028]** The computer-executable instructions (and the manipulated data) may be stored in the memory 104 of the computing system 100. Computing system 100 may also contain communication channels 108 that allow the computing system 100 to communicate with other computing systems over, for example, network 110.

10 **[0029]** While not all computing systems require a user interface, in some embodiments, the computing system 100 includes a user interface 112 for use in interfacing with a user. The user interface 112 may include output mechanisms 112A as well as input mechanisms 112B. The principles described herein are not limited to the precise output mechanisms 112A or input mechanisms 112B as such will depend on the  
15 nature of the device. However, output mechanisms 112A might include, for instance, speakers, displays, tactile output, holograms, virtual reality, and so forth. Examples of input mechanisms 112B might include, for instance, microphones, touchscreens, holograms, virtual reality, cameras, keyboards, mouse or other pointer input, sensors of any type, and so forth.

20 **[0030]** Embodiments described herein may comprise or utilize a special purpose or general-purpose computing system including computer hardware, such as, for example, one or more processors and system memory, as discussed in greater detail below. Embodiments described herein also include physical and other computer-readable media for carrying or storing computer-executable instructions and/or data structures. Such  
25 computer-readable media can be any available media that can be accessed by a general purpose or special purpose computing system. Computer-readable media that store computer-executable instructions are physical storage media. Computer-readable media that carry computer-executable instructions are transmission media. Thus, by way of example, and not limitation, embodiments can comprise at least two distinctly different  
30 kinds of computer-readable media: storage media and transmission media.

**[0031]** Computer-readable storage media includes RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other physical and tangible storage medium which can be used to store desired



program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computing system.

**[0032]** A “network” is defined as one or more data links that enable the transport of electronic data between computing systems and/or modules and/or other electronic devices. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computing system, the computing system properly views the connection as a transmission medium. Transmissions media can include a network and/or data links which can be used to carry desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computing system. Combinations of the above should also be included within the scope of computer-readable media.

**[0033]** Further, upon reaching various computing system components, program code means in the form of computer-executable instructions or data structures can be transferred automatically from transmission media to storage media (or vice versa). For example, computer-executable instructions or data structures received over a network or data link can be buffered in RAM within a network interface module (e.g., a “NIC”), and then eventually transferred to computing system RAM and/or to less volatile storage media at a computing system. Thus, it should be understood that readable media can be included in computing system components that also (or even primarily) utilize transmission media.

**[0034]** Computer-executable instructions comprise, for example, instructions and data which, when executed at a processor, cause a general purpose computing system, special purpose computing system, or special purpose processing device to perform a certain function or group of functions. Alternatively, or in addition, the computer-executable instructions may configure the computing system to perform a certain function or group of functions. The computer executable instructions may be, for example, binaries or even instructions that undergo some translation (such as compilation) before direct execution by the processors, such as intermediate format instructions such as assembly language, or even source code.

**[0035]** Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computing system configurations, including, personal computers, desktop computers, laptop computers, message processors, hand-held devices, multi-processor systems, microprocessor-based or programmable

consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, routers, switches, datacenters, wearables (such as glasses or watches) and the like. The invention may also be practiced in distributed system environments where local and remote computing systems, which are linked (either by  
5    hardwired data links, wireless data links, or by a combination of hardwired and wireless data links) through a network, both perform tasks. In a distributed system environment, program modules may be located in both local and remote memory storage devices.

**[0036]**   Those skilled in the art will also appreciate that the invention may be practiced in a cloud computing environment. Cloud computing environments may be distributed,  
10    although this is not required. When distributed, cloud computing environments may be distributed internationally within an organization and/or have components possessed across multiple organizations. In this description and the following claims, “cloud computing” is defined as a model for enabling on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and  
15    services). The definition of “cloud computing” is not limited to any of the other numerous advantages that can be obtained from such a model when properly deployed.

**[0037]**   For instance, cloud computing is currently employed in the marketplace so as to offer ubiquitous and convenient on-demand access to the shared pool of configurable computing resources. Furthermore, the shared pool of configurable computing resources  
20    can be rapidly provisioned via virtualization and released with low management effort or service provider interaction, and then scaled accordingly.

**[0038]**   A cloud computing model can be composed of various characteristics such as on-demand self-service, broad network access, resource pooling, rapid elasticity, measured service, and so forth. A cloud computing model may also come in the form of various  
25    service models such as, for example, Software as a Service (“SaaS”), Platform as a Service (“PaaS”), and Infrastructure as a Service (“IaaS”). The cloud computing model may also be deployed using different deployment models such as private cloud, community cloud, public cloud, hybrid cloud, and so forth. In this description and in the claims, a “cloud computing environment” is an environment in which cloud computing is employed.

**[0039]**   Figure 2 illustrates an environment 200 in which the principles described herein may operate. The environment 200 includes a physical space 201 that includes multiple physical entities 210, which may be any extant object, person, or thing that emits or reflects physical signals (such as electromagnetic radiation or acoustics) that has a pattern that may be used to potentially identify one or more physical features (also called  
30

herein states) of the respective object, person, or thing. An example of such potentially identifying electromagnetic radiation is visible light that has a light pattern (e.g., a still image or video) from which characteristics of visible entities may be ascertained. Such light pattern may be any temporal, spatial, or even higher-dimensional space. An example of such acoustics may be the voice of a human being, the sound of an object in normal operation or undergoing an activity or event, or a reflected acoustic echo.

**[0040]** The environment 200 also includes sensors 220 that receive physical signals from the physical entities 210. The sensors need not, of course, pick up every physical signal that the physical entity emits or reflects. For instance, a visible light camera (still or video) is capable of receiving electromagnetic radiation in the form of visible light and converting such signals into processable form, but cannot pick up all electromagnetic radiation of any frequency since cameras all have a finite dynamic range. Acoustic sensors likewise have limited dynamic range designed for certain frequency ranges. In any case, the sensors 220 provide (as represented by arrow 229) resulting sensor signals to a recognition component 230.

**[0041]** The recognition component 230 at least estimates (e.g., estimates or recognizes) one or more features of the physical entities 210 within the location based on patterns detected in the received sensor signals. The recognition component 230 may also generate a confidence level associated with the “at least an estimation” of a feature of the physical entity. If that confidence level is less than 100%, then the “at least an estimation” is just an estimation. If that confidence level is 100%, then the “at least an estimation” is really more than an estimation – it is a recognition. In the remainder of this description and in the claims, a feature that is “at least estimated” will also be referred to as a “sensed” feature to promote clarity. This is consistent with the ordinary usage of the term “sense” since a feature that is “sensed” is not always present with complete certainty. The recognition component 230 may employ deep learning (Deep Neural Network-based and reinforcement-based learning mechanisms) and machine learning algorithms to learn from experience what objects or people that are within an image, thereby improving accuracy of recognition over time.

**[0042]** The recognition component 230 provides (as represented by arrow 239) the sensed features into a sensed feature store 240, which can store the sensed features (and associated confidence levels) for each physical entity within the location 201, whether the physical entity is within the physical space for a short time, a long time, or permanently. The computation component 250 may then perform a variety of queries and/or

computations on the sensed feature data provided in sensed feature store 240. The queries and/or computations may be enabled by interactions (represented by arrow 249) between the computation component 250 and the sensed feature store 240.

5 [0043] In some embodiments, when the recognition component 230 senses a sensed feature of a physical entity within the location 201 using sensor signal(s) provided by a sensor, the sensor signals are also provided to a store, such as the sensed feature store. For instance, in Figure 2, the sensed feature store 240 is illustrated as including sensed features 241 as well as the corresponding sensor signals 242 that represent the evidence of the sense features.

10 [0044] For at least one (and preferably many) of the sensed features for at least one of the sensed plurality of entities, at least one signal segment is computer-associated with the sensed feature such that computer-navigation to the sensed feature also allows for computer-navigation to the signal segment. The association of the sensed signal with the associated signal segment may be performed continuously, thus resulting in an expanding graph, and an expanding collection of signal segments. That said, as described further below, garbage collection processes may be used to clean up sensed features and/or signal segments that are outdated or no longer of interest.

15 [0045] The signal segment may include multiple pieces of metadata such as, for instance, an identification of the sensor or sensors that generated the signal segment. The signal segment need not include all of the signals that were generated by that sensor, and for brevity, may perhaps include only those portions of the signal that were used to sense the sensed feature of the particular physical entity. In that case, the metadata may include a description of the portion of the original signal segment that was stored.

20 [0046] The sensed signal may be any type of signal that is generated by a sensor. Examples include video, image, and audio signals. However, the variety of signals is not limited to those that can be sensed by a human being. For instance, the signal segment might represented a transformed version of the signal generated by the sensor to allow for human observations of better human focus. Such transformations might include filtering, such a filtering based on frequency, or quantization. Such transformation might also include amplification, frequency shifting, speed adjustment, magnifying, amplitude adjustment, and so forth.

30 [0047] In order to allow for reduction in storage requirements as well as proper focus on the signal of interest, perhaps only a portion of the signal segment is stored. For instance, if a video signal, perhaps only a portion of the frames of the video are stored.

Furthermore, for any given image, perhaps only the relevant portion of the frame is stored. Likewise, if the sensor signal was an image, perhaps only the relevant portion of the image is stored. The recognition service that uses the signal segment to sense a feature is aware of which portion of the signal segment that was used to sense a feature. Accordingly, a  
5 recognition service can specifically carve out the relevant portion of the signal for any given sensed feature.

**[0048]** The computation component 250 may also have a security component 251 that may determine access to data with the sensed feature store 240. For instance, the security component 251 may control which users may access the sensed feature data 241 and/or the  
10 sensor signals 242. Furthermore, the security component 251 may even control which of the sensed feature data that computations are performed over, and/or which user are authorized to perform what type of computations or queries. Thus, security is effectively achieved. More regarding this security will be described below with respect to Figure 6.

**[0049]** Since the sensed feature data represents the sensed features of the physical  
15 entities within the physical space 201 over time, complex computing may be performed on the physical entities within the physical space 201. As will be described below, for a user, it is as though the very environment itself is filled with helpful computing power that is getting ready for any computing query or computation regarding that physical space. This will be referred to hereinafter also as “ambient computing”.

**[0050]** Furthermore, whenever a sensed feature is of interest, the evidence supporting  
20 that recognition components sensing of that feature may be reconstructed. For instance, the computing component 240 might provide video evidence of when a particular physical entity first entered a particular location. If multiple sensors generated sensor signals that were used by the recognition component to sense that feature, then the sensor signals for  
25 any individual sensor or combination of sensors may be reconstructed and evaluated. Thus, for instance, the video evidence of the physical entity first entering a particular location may be reviewed from different angles.

**[0051]** The physical space 201 is illustrated in Figure 2 and is intended just to be an abstract representation of any physical space that has sensors in it. There are infinite  
30 examples of such physical spaces, but examples include a room, a house, a neighborhood, a factory, a stadium, a building, a floor, an office, a car, an airplane, a spacecraft, a Petri dish, a pipe or tube, the atmosphere, underground spaces, caves, land, combinations and/or portions thereof. The physical space 201 may be the entirety of the observable universe or any portion thereof so long as there are sensors capable of receiving signals emitted from,

affected by (e.g., diffraction, frequency shifting, echoes, etc.), and/or reflected from the physical entities within the location.

5 [0052] The physical entities 210 within the physical space 201 are illustrated as including four physical entities 211, 212, 213 and 214 by way of example only. The ellipses 215 represent that there may be any number and variety of physical entities having features that are being sensed based on data from the sensors 220. The ellipses 215 also represent that physical entities may exit and enter the location 201. Thus, the number and identity of physical entities within the location 201 may change over time.

10 [0053] The position of the physical entities may also vary over time. Though the position of the physical entities is shown in the upper portion of the physical space 201 in Figure 2, this is simply for purpose of clear labelling. The principles described herein are not dependent on any particular physical entity occupying any particular physical position within the physical space 201.

15 [0054] Lastly, for convention only and to distinguish physical entities 210 from the sensors 220, the physical entities 210 are illustrated as triangles and the sensors 220 are illustrated as circles. The physical entities 210 and the sensors 220 may, of course, have any physical shape or size. Physical entities typically are not triangular in shape, and sensors are typically not circular in shape. Furthermore, sensors 220 may observe physical entities within a physical space 201 without regard for whether or not those sensors 220 are physically located within that physical space 201.

20 [0055] The sensors 220 within the physical space 201 are illustrated as including two sensors 221 and 222 by way of example only. The ellipses 223 represent that there may be any number and variety of sensors that are capable of receiving signals emitted, affected (e.g., via diffraction, frequency shifting, echoes, etc.) and/or reflected by the physical entities within the physical space. The number and capability of operable sensors may change over time as sensors within the physical space are added, removed, upgrade, broken, replaced, and so forth.

25 [0056] Figure 3 illustrates a flowchart of a method 300 for tracking physical entities within a physical space. Since the method 300 may be performed to track the physical entities 210 within the physical space 201 of Figure 2, the method 300 of Figure 3 will now be described with frequent reference to the environment 200 of Figure 2. Also, Figure 4 illustrates an entity tracking data structure 400 that may be used to assist in performing the method 300, and which may be used to later perform queries on the tracked physical entities, and perhaps also to access and review the sensor signals

associated with the tracked physical entities. Furthermore, the entity tracking data structure 400 may be stored in the sensed feature store 240 of Figure 4 (which is represented as sensed feature data 241). Accordingly, the method 300 of Figure 3 will also be described with frequent reference to the entity tracking data structure 400 of Figure

5 4.

**[0057]** In order to assist with tracking, a space-time data structure for the physical space is set up (act 301). This may be a distributed data structure or a non-distributed data structure. Figure 4 illustrates an example of an entity tracking data structure 400 that includes a space-time data structure 401. This entity tracking data structure 400 may be included within the sensed feature store 240 of Figure 2 as sensed feature data 241. While the principles described herein are described with respect to tracking physical entities, and their sensed features and activities, the principles described herein may operate to tracking physical entities (and their sensed features and activities) within more than one location. In that case, perhaps the space-time data structure 401 is not the root node in the tree represented by the entity tracking data structure 400 (as symbolized by the ellipses 402A and 402B). Rather there may be multiple space-time data structures that may be interconnected via a common root node.

**[0058]** Then, returning to Figure 3, the content of box 310A may be performed for each of multiple physical entities (e.g., physical entities 210) that are at least temporarily within a physical space (e.g., physical space 201). Furthermore, the content of box 310B is illustrated as being nested within box 310A, and represents that its content may be performed at each of multiple times for a given physical entity. By performing the method 300, a complex entity tracking data structure 400 may be created and grown, to thereby record the sensed features of physical entities that are one or more times within the location. Furthermore, the entity tracking data structure 400 may potentially also be used to access the sensed signals that resulted in certain sensed features (or feature changes) being recognized.

**[0059]** For a particular physical entity in the location at a particular time, a physical entity is sensed by one or more sensors (act 311). In other words, one or more physical signals emitted from, affected by (e.g., via diffraction, frequency shifting, echoes, etc.), and/or reflected from the physical entity is received by one or more of the sensors. Referring to Figure 1, suppose that physical entity 211 has one or more features that are sensed by both sensors 221 and 222 at a particular time.

[0060] One aspect of security may enter at this point. The recognition component 230 may have a security component 231 that, according to particular settings, may refuse to record sensed features associated with particular physical entities, sensed features of a particular type, and/or that were sensed from sensor signals generated at particular time, or combinations thereof. For instance, perhaps the recognition component 230 will not record sensed features of any people that are within the location. As a more fine-grained examples, perhaps the recognition component 230 will not record sensed features of a set of people, where those sensed features relate to an identity or gender of the person, and where those sensed features resulted from sensor signals that were generated at particular time frames. More regarding this security will again be described below with respect to Figure 6.

[0061] If permitted, an at least approximation of that particular time at which the physical entity was sensed is represented within an entity data structure that corresponds to the physical entity and this is computing-associated with the space-time data structure (act 312). For instance, referring to Figure 4, the entity data structure 410A may correspond to the physical entity 211 and is computing-associated (as represented by line 430A) with the space-time data structure 401. In this description and in the claims, one node of a data structure is “computing-associated” with another node of a data structure if a computing system is, by whatever means, able to detect an association between the two nodes. For instance, the use of pointers is one mechanism for computing-association. A node of a data structure may also be computing-associated by being included within the other node of the data structure, and by any other mechanism recognized by a computing system as being an association.

[0062] The time data 411 represents an at least approximation of the time that the physical entity was sensed (at least at this time iteration of the content of box 310B) within the entity data structure 410A. The time may be a real time (e.g., expressed with respect to an atomic clock), or may be an artificial time. For instance, the artificial time may be a time that is offset from real-time and/or expressed in a different manner than real time (e.g., number of seconds or minutes since the last turn of the millennium). The artificial time may also be a logical time, such as a time that is expressed by a monotonically increasing number that increments at each sensing.

[0063] Also, based on the sensing of the particular physical entity at the particular time (at act 311), the environment senses at least one physical feature (and perhaps multiple) of the particular physical entity in which the particular physical entity exists at



the particular time (act 313). For instance, referring to Figure 2, the recognition component 230 may sense at least one physical feature of the physical entity 211 based on the signals received from the sensors 221 and 222 (e.g., as represented by arrow 229).

**[0064]** The sensed at least one physical feature of the particular physical entity is then represented in the entity data structure (act 314) in a manner computing-associated with the at least approximation of the particular time. For instance, in Figure 2, the sensed feature data is provided (as represented by arrow 239) to the sensed feature store 240. In some embodiments, this sensed feature data may be provided along with the at least approximation of the particular time so as to modify the entity tracking data structure 400 in substantially one act. In other words, act 312 and act 314 may be performed at substantially the same time to reduce write operations into the sensed feature store 240.

**[0065]** Furthermore, if permitted, the sensor signal(s) that the recognition component relied upon to sense the sensed feature are recorded in a manner that is computer-associated with the sensed feature (act 315). For instance, the sensed feature that is in the sensed feature data 241 (e.g., in the space-time data structure 401) may be computing-associated with such sensor signal(s) stored in the sensed signal data 242.

**[0066]** Referring to Figure 4, the first entity data structure now has sensed feature data 421 that is computing-associated with time 411. In this example, the sensed feature data 421 includes two sensed physical features 421A and 421B of the physical entity. However, the ellipses 421C represents that there may be any number of sensed features of the physical entity that is stored as part of the sensed feature data 421 within the entity data structure 401. For instance, there may be a single sensed feature, or innumerable sensed features, or any number in-between for any given physical entity as detected at any particular time.

**[0067]** In some cases, the sensed feature may be associated with other features. For instance, if the physical entity is a person, the feature might be a name of the person. That specifically identified person might have known characteristics based on features not represented within the entity data structure. For instance, the person might have a certain rank or position within an organization, have certain training, be a certain height, and so forth. The entity data structure may be extended by, when a particular feature is sensed (e.g., a name), pointing to additional features of that physical entity (e.g., rank, position, training, height) so as to even further extend the richness of querying and/or other computation on the data structure.

[0068] The sensed feature data may also have confidence levels associated with each sensed feature that represents an estimated probability that the physical entity really has the sensed feature at the particular time 410A. In this example, confidence level 421a is associated with sensed feature 421A and represents a confidence that the physical entity  
5 211 really has the sensed feature 421A. Likewise, confidence level 421b is associated with sensed feature 421B and represents a confidence that the physical entity 211 really has the sensed feature 421B. The ellipses 421c again represents that there may be confidence levels expressed for any number of physical features. Furthermore, there may be some physical features for which there is no confidence level expressed (e.g., in the  
10 case where there is certainty or in case where it is not important or desirable to measure confidence of a sensed physical feature).

[0069] The sensed feature data may also have computing-association (e.g., a pointer) to the sensor signal(s) that were used by the recognition component to sense the sense feature of that confidence level. For instance, in Figure 4, sensor signal(s) 421Aa is  
15 computing-associated with sensed feature 421A and represents the sensor signal(s) that were used to sense the sensed feature 421A at the time 411. Likewise, sensor signal(s) 421Bb is computing-associated with sensed feature 421B and represents the sensor signal(s) that were used to sense the sensed feature 421B at the time 411. The ellipses 421Cc again represents that there may be computing-associations of any number of  
20 physical features.

[0070] The security component 231 of the recognition component 230 may also exercise security in deciding whether or not to record sensor signal(s) that were used to sense particular features at particular times. Thus, the security component 231 may exercise security in 1) determining whether to record that particular features were sensed,  
25 2) determining whether to record features associated with particular physical entities, 3) determining whether to record features sensed at particular times, 4) determining whether to record the sensor signal(s), and if so which signals, to record as evidence of a sensed feature, and so forth.

[0071] As an example, suppose that the location being tracked is a room. Now  
30 suppose that an image sensor (e.g., a camera) senses something within the room. An example sensed feature is that the “thing” is a human being. Another example sensed feature is that the “thing” is a particular named person. There might be a confidence level of 100 percent that the “thing” is a person, but only a 20 percent confidence level that the person is a specific identified person. In this case, the sensed feature set includes one

feature that is a more specific type of another feature. Furthermore, the image data from the camera may be pointed to by the record of the sensed feature of the particular physical entity at the particular time.

[0072] Another example feature is that the physical entity simply exists within the location, or at a particular position within the location. Another example is that this is the first appearance of the physical entity since a particular time (e.g., in recent times, or even ever). Another example of features is that the item is inanimate (e.g., with 99 percent certainty), a tool (e.g., with 80 percent certainty), and a hammer (e.g., with 60 percent certainty). Another example feature is that the physical entity is no longer present (e.g., is absent) from the location, or has a particular pose, is oriented in a certain way, or has a positional relationship with another physical entity within the location (e.g., “on the table” or “sitting in chair #5”).

[0073] In any case, the number and types of features that can be sensed from the number and types of physical entities within any location is innumerable. Also, as previously mentioned, as represented by box 310B, the acts within box 310B may potentially be performed multiple times for any given physical entity. For instance, physical entity 211 may be against detected by one or both of sensors 221 and 222. Referring to Figure 4, this detection results in the time of the next detection (or is approximation) to be represented within the entity data structure 410. For instance, time 412 is also represented within the entity data structure. Furthermore, sensed features 422 (e.g., including perhaps sensed feature 422A and 422B -with ellipses 422C again representing flexibility) are computing-associated with the second time 412. Furthermore, those sensed features may also have associated confidence levels (e.g., 422a, 422b, ellipses 422c). Likewise, those sensed features may also have associated sensor signals (e.g., 422Aa, 422Bb, ellipses 422Cc).

[0074] The sensed features sensed at the second time may be the same as or different than the sensed features sensed at the first time. The confidence levels may change over time. As an example, suppose a human being is detected at time #1 at one side of a large room via an image with 90 percent confidence, and that the human being is specifically sensed as being John Doe with 30 percent confidence. Now, at time #2 that is 0.1 seconds later, John Doe is sensed 50 feet away at another part of the room with 100 percent confidence, and there remains a human being at the same location where John Doe was speculated to be at time 1. Since human beings do not travel 50 feet in a tenth of a second (at least in an office setting), it can now be concluded that the human being detected at

time 1 is not John Doe at all. So that confidence for time #1 that the human being is John Doe is reduced to zero.

[0075] Returning to Figure 2, the ellipses 413 and 423 represent that there is no limit to the number of times that a physical entity may be detected within the location. As subsequent detections are made, more may be learned about the physical entity, and thus sensed features may be added (or removed) as appropriate, with corresponding adjustments to confidence levels for each sensed feature.

[0076] Now moving outside of box 310B, but remaining within box 310A, for any given physical entity, feature changes in the particular entity may be sensed (act 322) based on comparison (act 321) of the sensed feature(s) of the particular physical entity at different times. This sensed changes may be performed by the recognition component 230 or the computation component 250. If desired, those sensed changes may also be recorded (act 323). For instance, the sensed changes may be recorded in the entity data structure 410A in a manner that is, or perhaps is not, computing-associated with a particular time. Sensor signals evidencing the feature change may be reconstructed using the sensor signals that evidenced the sensed feature at each time.

[0077] For instance, based on a sensed feature at a first time being a presence of the physical entity within the location, and based on a second feature at a second time being an absence of the physical entity within the location, it can be concluded that the physical entity has exited the physical space. On the contrary, based on a sensed feature at a first time being an absence of the physical entity from the location, and a second feature at a second time being a presence of the physical entity within the location, it can be concluded that the physical entity has entered the location. In some case, perhaps absence from a physical space is not looked for in a physical entity until the physical entity is first detected as being present in the physical space.

[0078] Now referring to the box 310A, this tracking of feature(s) of physical entities may be performed for multiple entities over time. For instance, the content of box 310A may be performed for each of physical entities 211, 212, 213 or 214 within the physical space 201 or for other physical entities that enter or exit the physical space 201. Referring to Figure 4, the space-time data structure 401 also is computing-associated (as represented by lines 430B, 430C, and 430D) with a second entity data structure 410B (perhaps associated with the second physical entity 212 of Figure 2), a third entity data structure 410C (perhaps associated with the third physical entity 213 of Figure 2); and a fourth

entity data structure 410D (perhaps associated with the fourth physical entity 214 of Figure 2).

**[0079]** The space-time data structure 401 may also include one or more triggers that define conditions and actions. When the conditions are met, corresponding actions are to occur. The triggers may be stored at any location in the space-time data structure. For instance, if the conditions are/or actions are with respect to a particular entity data structure, the trigger may be stored in the corresponding entity data structure. If the conditions and/or actions are with respect to a particular feature of a particular entity data structure, the trigger may be stored in the corresponding feature data structure.

**[0080]** The ellipses 410E represent that the number of entity data structures may change. For instance, if tracking data is kept forever with respect to physical entities that are ever within the physical space, then additional entity data structures may be added each time a new physical entity is detected within the location, and any given entity data structure may be augmented each time a physical entity is detected within the physical space. Recall, however, that garbage collection may be performed (e.g., by clean-up component 260) to keep the entity tracking data structure 400 from growing too large to be properly edited, stored and/or navigated.

**[0081]** Outside of the box 310A, physical relationships between different physical entities may be sensed (act 332) based on comparison of the associated entities data structures (act 331). Those physical relationships may likewise be recorded in the entity tracking data structure 401 (act 333) perhaps within the associated entity data structures that have the sensed physical relationships, and/or perhaps associated with the time that the physical entities are sensed as having the relationship. For instance, by analysis of the entity data structures for different physical entities through time, it might be determined that at a particular time, that a physical entity may be hidden behind another physical entity, or that a physical entity may be obscuring the sensing of another physical entity, or that two physical entities have been joined or that a physical entity has been detached to create multiple physical entities. Sensor signals evidencing the physical entity relationship may be reconstructed using the sensor signals that evidenced the sensed feature at the appropriate time and for each physical entity.

**[0082]** The feature data store 240 may now be used as a powerful store upon which to compute complex functions and queries over representations of physical entities over time within a physical space. Such computation and querying may be performed by the computation component 250. This enables enumerable numbers of helpful embodiments,

and in fact introduces an entirely new form of computing referred to herein as “ambient computing”. Within the physical space that has sensors, it is as though the very air itself can be used to compute and sense state about the physical world. It is as though a crystal ball has now been created for that physical space from which it is possible to query and/or compute many things about that location, and its history.

**[0083]** As an example, a user may now query whether an object is right now in a physical space, or where an object was at a particular time within the physical space. The user might also query which person having particular features (e.g., rank or position within a company) is near that object right now, and communicate with that person to bring the object to the user. The user might query as to relationships between physical entities. For instance, the user might query who has possession of an object. The user might query as to the state of an object, whether it is hidden, and what other object is obscuring view of the object. The user might query when a physical entity first appeared within the physical space, when they exited, and so forth. The user might also query when the lights were turned off, when the system became certain of one or more features of a physical entity. The user might also search on feature(s) of an object. The user might also query on activities that have occurred within the location. A user might compute the mean time that a physical entity of a particular type is within the location, anticipate where a physical entity will be at some future time, and so forth. Accordingly, rich computing and querying may be performed on a physical space that has sensors.

**[0084]** As previously mentioned, the computer-navigable graph may have signal segments associated with sensed features. Figure 5 illustrates a flowchart of a method 500 for efficiently rendering signal segments of interest. First, the computing system navigates the navigable graph of sensed features to reach a particular sensed feature (act 501). For instance, this navigation may be performed automatic or in response to user input. The navigation may be the result of a calculation, or may simply involve identifying the sensed feature of interest. As another example, the navigation may be the result of a user query. In some embodiments, a calculation or query may result in multiple sensed features being navigated to. As an example, suppose that the computing system navigates to sensed feature 222A in Figure 2.

**[0085]** The computing system then navigates to the sensed signal computer-associated with the particular sensed feature (act 502) using the computer-association between the particular sensed feature and the associated sensor signal. For instance, in Figure 2, with

the sensed feature being sensed feature 222A, the computer-association is used to navigate to the signal segment 222Aa.

[0086] Finally, the signal segment may then be rendered (act 503) on an appropriate output device. For instance, if the computing system is the computing system 100 of Figure 1, the appropriate output device might be one or more of output mechanisms 112A. For instance, audio signals may be rendered using speakers, and visual data may be rendered using a display. After navigating to the sensed signal(s), multiple things could happen. The user might play a particular signal segment, or perhaps choose from multiple signal segments that contributed to the feature. A view could be synthesized from the multiple signal segments.

[0087] With computing being performed on the physical world, a new type of ambient computation is enabled. It is as though computers are available in the very ambient environment, embodied within the air itself, and able to perform computations on physical entities that were at any point in contact with that air. In the workplace, productivity may be greatly improved using this ambient computing. For instance, a user may quickly find a misplaced tool, or be able to communicate with a peer close to the tool so that the user can ask that peer to grab that tool and bring it to the user. Furthermore, in addition to ambient computing, human beings may review the sensor signal(s) that were used to sense features of interest for particular physical entities of interest, at particular times of interest. However, the number of scenarios for improving physical productivity by due to responsible use of ambient computing is limitless.

[0088] Now that the principles of ambient computing have been described with respect to Figures 2 through 5, security mechanisms that may be performed in the context of such ambient computing will be described with respect to Figure 6. Figure 6 illustrates a flowchart of a method 600 for controlling creation of or access to information sensed by one or more sensors in a physical space. The method includes creating (act 601) a computer-navigable graph of features of sensed physical entities sensed in a physical space over time. The principles described herein are not limited to the precise structure of such a computer-navigable graph. An example structure and its creation have been described with respect to Figures 2 through 4.

[0089] The method 600 also includes restricting creation of or access to nodes of the computer-navigable graph based on one or more criteria (act 602). Thus, security is imposed upon the computer-navigable graph. The arrows 603 and 604 represent that the process of creating the graph and restrict creation/access to its nodes may be a continual

process. The graph may be continuously have nodes added to (and perhaps removed from) the graph. Furthermore, restrictions of creation may be considered whenever there is a possibility of creation of a node. Restrictions of access may be decided when a node of the graph is created, or at any point thereafter. Examples of restrictions might include, for instance, a prospective identity of a sensed physical entity, a sensed feature of a sensed physical entity, and so forth.

**[0090]** In determining whether access to a node of a computer-navigable graph is authorized, there may be access criteria for each node. Such access criteria may be explicit or implicit. That is, if there is no access criteria explicit for the node that is to be accessed, then perhaps a default set of access criteria may apply. The access criteria for any given node may be organized in any manner. For instance, in one embodiment, the access criteria for a node may be stored with the node in the computer-navigable graph.

**[0091]** The access restrictions might also include restrictions based on a type of access requested. For instance, a computational access means that node is not directly accessed, but is used in a computation. Direct access to read the content of a node may be restricted, whilst computational access that does not report the exact contents of the node may be allowed.

**[0092]** Access restrictions may also be based on the type of node accessed. For instance, there may be a restriction in access to the particular entity data structure node of the computer-navigable graph. For instance, if that particular entity data structure node represents detections of a particular person in the physical space, access might be denied. There may also be restrictions in access to particular signal segment nodes of the computer-navigable graph. As an example, perhaps one may be able to determine that a person was in a location at a given time, but not be able to review video recordings of that person at that location. Access restrictions may also be based on who is the requestor of access.

**[0093]** In determining whether to restrict creation of a particular sensed feature node of the computer-navigable graph, there may be a variety of criteria considered. For instance, there may be a restriction in creation of a particular signal segment node of a computer-navigable graph.

**[0094]** Figure 7 illustrates a recurring flow 700 showing that in addition to creating a computer-navigable graph of sensed features in the physical space (act 701), there may also be pruning of the computer-navigable graph (act 702). These acts may even occur simultaneously and continuously (as represented by the arrows 703 and 704) to thereby



keep the computer-navigable graph of sensed features at a manageable size. There has been significant description herein about how the computer-navigable graph may be created (represented as act 701).

**[0095]** Now, this description will focus on how the computer-navigable graph may be pruned to remove one or more nodes of the computer-navigable graph (act 702). Any node of the computer-navigable graph may be subject to removal. For instance, sensed features of a physical entity data structure may be removed for specific time or group of times. A sensed feature of a physical entity data structure may also be removed for all times. More than one sensed features of a physical entity data structure may be removed for any given time, or for any group of times. Furthermore, a physical entity data structure may be entirely removed in some cases.

**[0096]** The removal of a node may occur, for instance, when the physical graph represents something that is impossible given the laws of physics. For instance, a given object cannot be at two places at the same time, nor can that object travel significant distances in a short amount of time in an environment in which such travel is infeasible or impossible. Accordingly, if a physical entity is tracked with absolute certainty at one location, any physical entity data structure that represent with lesser confidence that the same physical entity is at an inconsistent location may be deleted.

**[0097]** The removal of a node may also occur when more confidence is obtained regarding a sensed feature of a physical entity. For instance, if a sensed feature of a physical entity within a location is determined with 100 percent certainty, then the certainty levels of that sensed feature of that physical entity may be updated to read 100 percent for all prior times also. Furthermore, sensed features that have been learned to not be applicable to a physical entity (i.e., the confidence level has reduced to zero or negligible), the sensed feature may be removed for that physical entity.

**[0098]** Furthermore, some information in the computer-navigable graph may simply be too stale to be useful. For instance, if a physical entity has not been observed in the physical space for a substantial period of time so as to make the prior recognition of the physical entity no longer relevant, then the entire physical entity data structure may be removed. Furthermore, detections of a physical entity that have become staled may be removed though the physical entity data structure remains to reflect more recent detections. Thus, cleansing (or pruning) of the computer-navigable graph may be performed via intrinsic analysis and/or via extrinsic information. This pruning intrinsically improves the quality of the information represented in the computer-navigable graph, by

removing information of lesser quality, and freeing up space for more relevant information to be stored.

**[0099]** Accordingly, the principles described herein allow for a computer-navigable graph of the physical world. The graph may be searchable and queriable thereby allowing  
5 for searching and querying and other computations to be performed on the real world. Security may further be imposed in such an environment. Finally, the graph may be kept to a manageable size through cleansing and pruning. Thus, a new paradigm in computing has been achieved.

**[00100]** The above-described computer-navigable graph of physical space enables a  
10 wide variety of applications and technical achievements. In particular, three of such achievements that will now be described are based on a physical graph that has signal segments that evidence state of physical entities, and thus the physical graph has a semantic understanding of what is happening within any given signal segment. In a first implementation, portions of signal segments may be shared, with semantic understanding  
15 taking part on which portions of the signal segment are extracted for sharing. In a second implementation, signal segments may automatically be narrated using semantic understanding of what is happening within a signal segment. In a third implementation, actors may be trained in a just-in-time fashion by providing representations of signal segments when the actor is to begin or has begun an activity.

**[00101]** Figure 8 illustrates a flowchart of a method 800 for sharing at least a portion of  
20 a signal segment. The signal segment might be, for instance, multiple signal segments that have captured the same physical entity. For instance, if the signal segment is a video signal segment, multiple video segments may have captured the same physical entity or entities from different perspectives and distances. If the signal is an audio signal segment,  
25 multiple audio segments may have been captured the selected physical entity or entities with different acoustic channels intervening between corresponding acoustic sensors and the selected physical entity or entities (or portions thereof). The signal segment(s) being shared may be a live signal segment that is capturing signals live from one or more physical entities within a location. Alternatively, the signal segment(s) being shared may  
30 be a recorded signal segment.

**[00102]** In accordance with the method 800, the system detects selection of one or more physical entities or portions thereof that is/are rendered within one or more signal segments (act 801). Thus, sharing may be initiated based on the semantic content of a signal segment. For instance, the selected physical entity or entities (or portion(s) thereof)

may be the target of work or a source of work. As an example, the user might select a target of work such as a physical whiteboard. Another example target of work might be a piece of equipment that is being repaired. Examples of sources of work might be, for instance, a person writing on a physical whiteboard, a dancer, a magician, a construction worker, and so forth.

**[00103]** The individual that selected the physical entity or entities (or portions thereof) for sharing may be a human being. In that case, the user might select the physical entity or entities (or portions thereof) in any manner intuitive to a human user. Examples of such input include gestures. For instance, the user might circle an area that encompasses the physical entity or entities (or portions thereof) within a portion of a video or image signal segment.

**[00104]** Alternatively, the selection may be made by a system. For instance, the system might select that the portion of signal segments that includes particular physical entity or entities (or portions thereof) be shared upon detection of a particular condition, and/or in accordance with policy. For instance, as described below with respect to Figure 10, the system might detect that a human actor is about to engage in a particular activity that requires training. The system might then select the physical entities or entities that are similar to a target of activity, or that include an individual as that individual has previously performed the activity, to share with the human actor. A narration of the activity may even be automatically generated and provide (as described with respect to Figure 9).

**[00105]** The system then extracts portion(s) of the signal segment(s) in which the selected physical entity or selected portion of the physical entity is rendered (act 802). For instance, the signal segment might be multiple video signal segments. The system might create a signal segment in which the point of view changes from one signal segment (generated by one sensor) to another signal segment (generated by another sensor) upon the occurrence of condition(s) that occur with respect to the selected physical entity or entities (or the selected portion there). For instance, suppose the selected physical entity is those portions of the whiteboard that an instructor is currently writing on. If the instructor's body was to obscure his own writing from the perspective of one sensor, another signal segment that captures the active portion of the whiteboard may be switched to automatically. The system may perform such switching (of live signal segments) or stitching (or recorded video segments) automatically.

**[00106]** The system then dispatches a representation of the signal segment(s) that encompasses the selected physical entity or entities (or portions there) to one or more

recipients (act 803). Such recipients may be human beings, components, robotics, or any other entity capable of using the shared signal segment portion(s).

[00107] In one embodiment, the signal segments represents a portion of a physical graph that includes representations of physical entities sensed within a physical space, along with signal segments that evidence state of the physical entities. An example of such a physical graph has been described above with respect to Figures 2 through 4 with respect to the computer-navigable graph 400. The system could also dispatch a portion of the physical graph that relates to the signal segment portion(s) that are shared, and/or perhaps may extract information from that corresponding portion of the physical graph to share along with (or as an alternative to) the sharing of the signal segment portion(s) themselves.

[00108] As previously mentioned, the representation of the shared portion could be an automatically generated narration. The semantic understanding of what physical activity and entities is being depicted in a signal segment also allows for automatic generation of a narration of the signal segment (whether live narration, or narration of a recorded video segment). Figure 9 illustrates a flowchart of a method 900 for automatically generating a narration of what is happening in a signal segment. As an example, automatic narration of a chess game, a football game, or the like, may be performed. Automatic narration of a work activity may also be performed.

[00109] To generate an automatic narration of a signal segment, the signal segment is accessed from the physical graph (act 901). Using the semantic understanding of what is depicted in the signal segment, the system then determines how physical entity or entities are acting in the signal segment (act 902). Again, such a semantic understanding may be gained from the physical graph portions that correspond to the signal segment.

[00110] Not everything rendered in a signal segment will be of relevance for narration. The system might determine what actions are salient taking into consideration any number of factors (or a balance thereof) including, for instance, whether an action is happening repeatedly, what has not changed, what is occurring constantly, what portion(s) of the signal segment have been shared, user instruction, and so forth. There could also be some training of the system and machine learning to know what actions are potentially relevant.

[00111] The system then automatically generates a narration of the actions in the signal segment using the determined actions of the one or more physical entities (act 903). As an example, if the signal segment were a video signal segment, the generated narration might be one or more of an audio narration, a diagrammatic narration, or the like. In the audio

narration, what is happening in the video signal segment might be spoken. In a diagrammatic narration, a diagram of what is happening in the video signal segment may be rendered with irrelevant material removed, and with relevant physical entities being visually emphasized. The relevant physical entities might be represented in simplified  
5 (perhaps cartoonish) form, with movement and actions potentially visually represented (e.g., by arrows).

**[00112]** If the signal segment were an audio signal segment, the generated narration might include a summarized audio narration that includes a simplified audio of the relevant matters that are heard within the signal segment. The generated narration might  
10 also be a diagrammatic narration. The narration might also show or describe to an intended recipient what to do with respect to the intended recipient's environment, since the physical graph might also have semantic understanding of the surroundings of the intended recipient.

**[00113]** Figure 10 illustrates a flowchart of a method 1000 for automatically training  
15 upon the occurrence of a condition with respect to that actor. The method 1000 is initiated upon detecting that the condition is met with respect to the actor a human actor is engaging in or is about to engage in a physical activity (act 1001). For instance, the condition might be that the actor is performing or is about to perform a physical activity, or that the actor has a certain physical status (e.g., presence within a room). In response,  
20 the system evaluates whether or not training is to be provided to the human actor for that physical activity (act 1002). For instance, the determination might be based on some training policy. As an example, training might be mandatory at least once for all employees of an organization with respect to that activity, and might be mandated to occur before the activity is performed. Training might be required on an annual basis. Training  
25 might be required depending on who the actor is (e.g., a new hire, or a safety officer).

**[00114]** Training might be offered when the system determines that the human actor is engaging in an activity improperly (e.g., failing to bend her knees when lifting a heavy object). The training may be tailored to how the human actor is performing the activity improperly. The training may be offered repeatedly depending on a learning rate of the  
30 human actor.

**[00115]** Training is then dispatched to the actor (act 1003). For instance, a robot or human may be dispatched to the actor to show the actor how to perform an activity. Alternatively or in addition, a representation of a signal segment may then be dispatched to the actor (act 1003), where the representation provides training to the human actor. For

instance, the representation might be a narration (such as that automatically generated by the method 900 of Figure 9). The shared representation might be a multi-sensor signal segment – such as a stitched video signal in which the point of view strategically changes upon the occurrence of one or more conditions with respect to a selected physical entity or selected portion of the physical entity, as described above with respect to Figure 9.

**[00116]** Accordingly, the principles described herein use an ambient computing environment, and the semantic understanding of what is happening in the real world, in order to provide significant technical advancements. The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

## CLAIMS

1. A computing system comprising:  
one or more processors;  
one or more computer-readable media having thereon computer-executable instructions that are structured such that, when executed by the one or more processors, cause the computing system to perform a method for automatically training an actor upon the occurrence of a condition with respect to the actor, the method comprising:  
detecting that a condition has occurred with respect to an actor;  
in response to the detection, determining that training is to be provided to the actor regarding the physical activity; and  
in response to determining, dispatching training to the actor.
2. The computing system in accordance with Claim 1, the condition comprising that the actor is engaging or is about to engage in a physical activity.
3. The computing system in accordance with Claim 1, the dispatching of training to the actor comprising:  
dispatching a representation of a signal segment to the actor, the representation providing training to the actor.
4. The computing system in accordance with Claim 3, the condition with respect to the actor being that the actor is performing the physical activity improperly, the representation of the signal segment being different depending on how the actor is performing the physical activity improperly.
5. The computing system in accordance with Claim 3, the signal segment comprising a multi-sensor signal segment.
6. The computing system in accordance with Claim 3, the signal segment comprising a video signal segment.
7. The computing system in accordance with Claim 3, a physical entity represented in the signal segment being a target of work similar to a target of work to be or being operated upon by the actor in the activity.
8. The computing system in accordance with Claim 3, a physical entity represented in the signal segment being someone that is performing the activity in the past, the signal segment being a multi-video signal segment in which the point of view changes to a different sensor upon the occurrence of one or more conditions with respect to the selected physical entity or selected portion of the physical entity.

9. The computing system in accordance with Claim 1, the condition with respect to the actor being that the actor is performing the physical activity improperly.

10. A method for automatically training an actor upon the occurrence of a condition with respect to the actor, the method comprising:

detecting that a condition has occurred with respect to an actor;

in response to the detection, determining that training is to be provided to the actor regarding the physical activity; and

in response to determining, dispatching training to the actor.




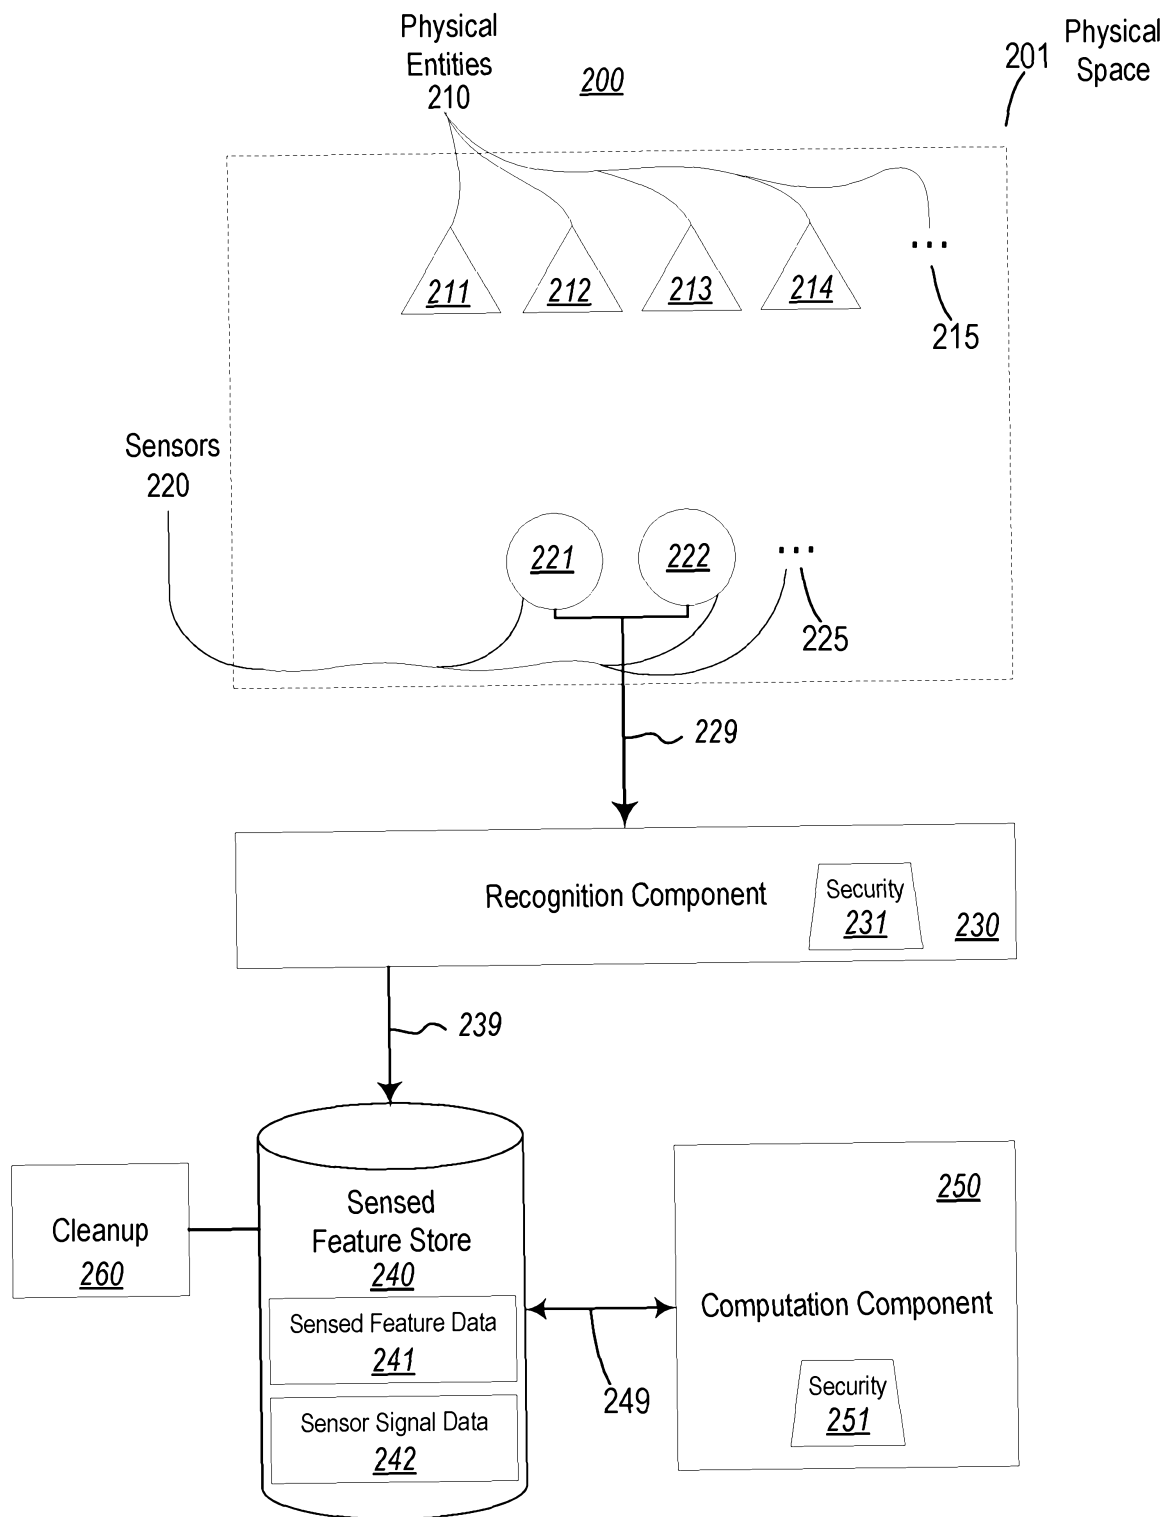
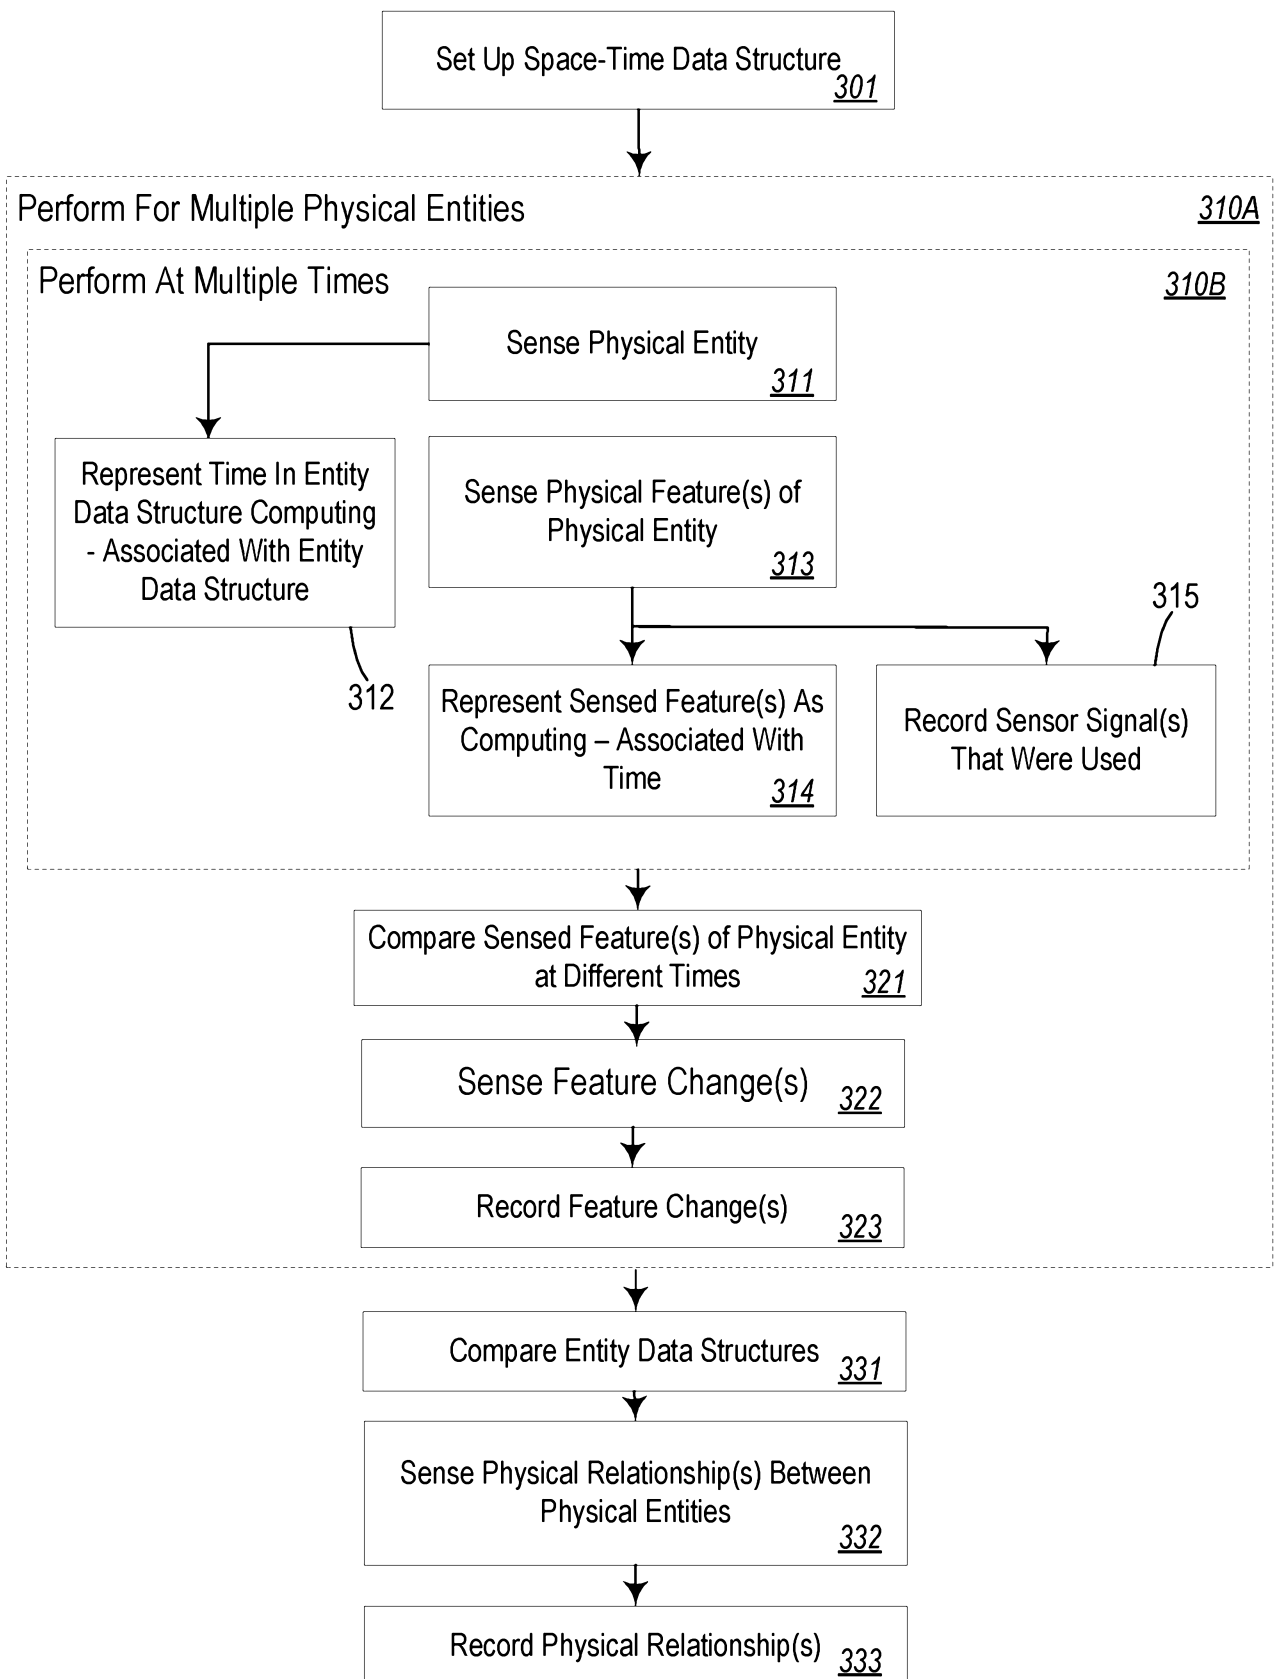


Figure 1

2/8

**Figure 2**

3/8

300**Figure 3**


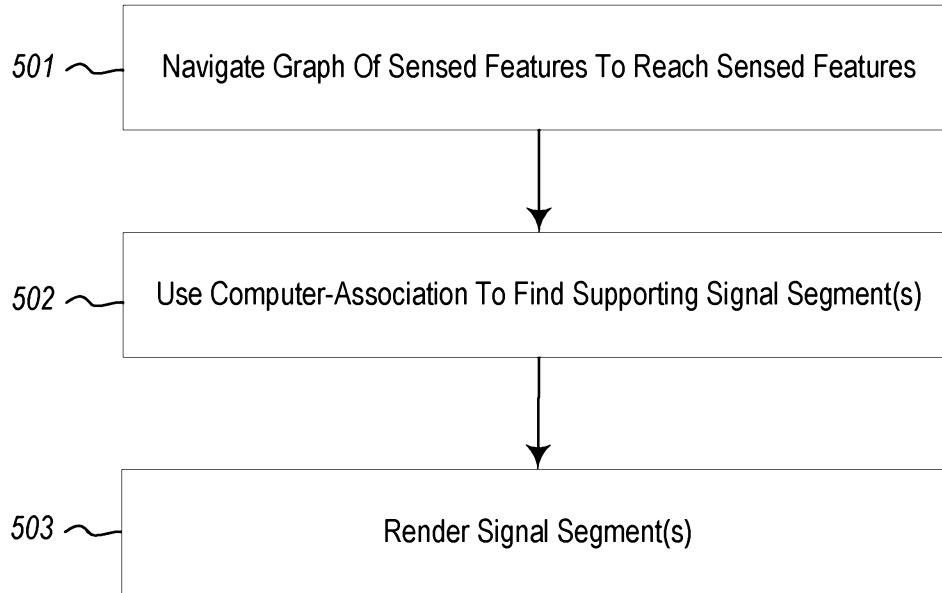
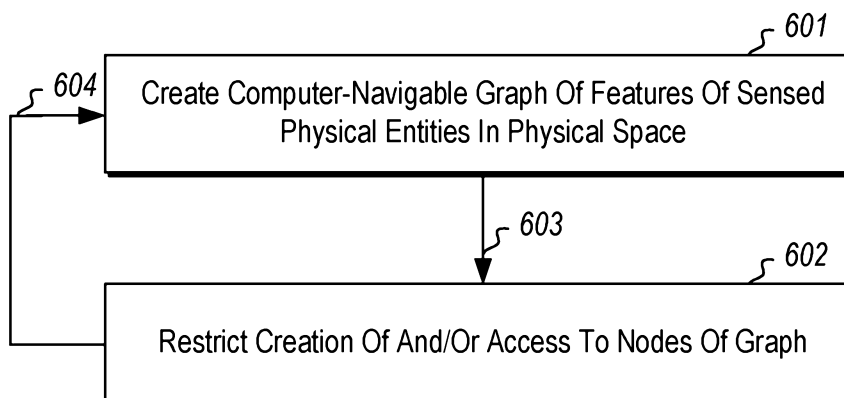
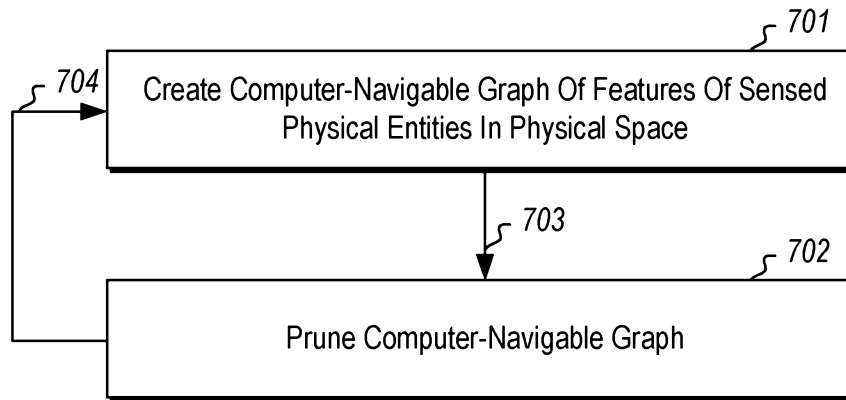
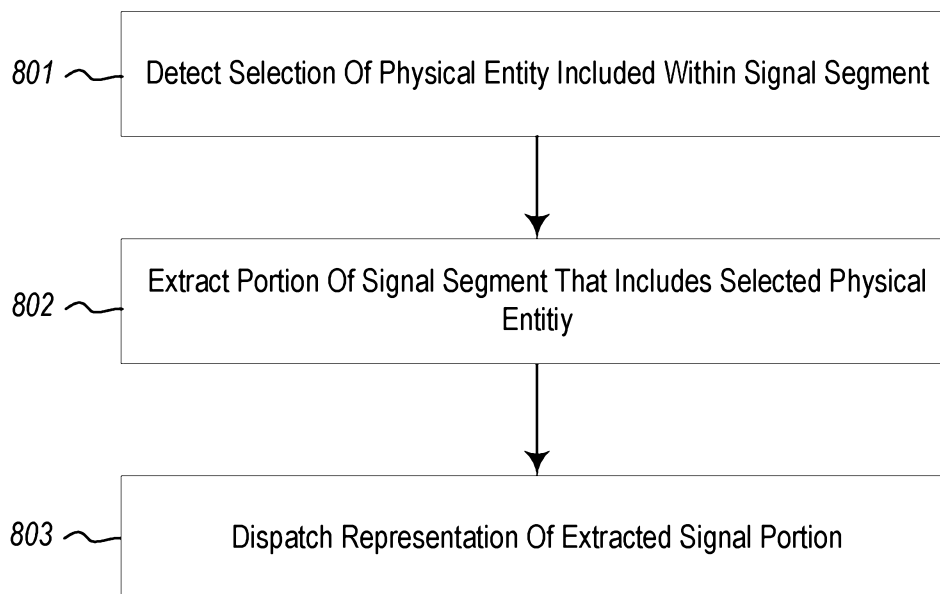


Figure 4

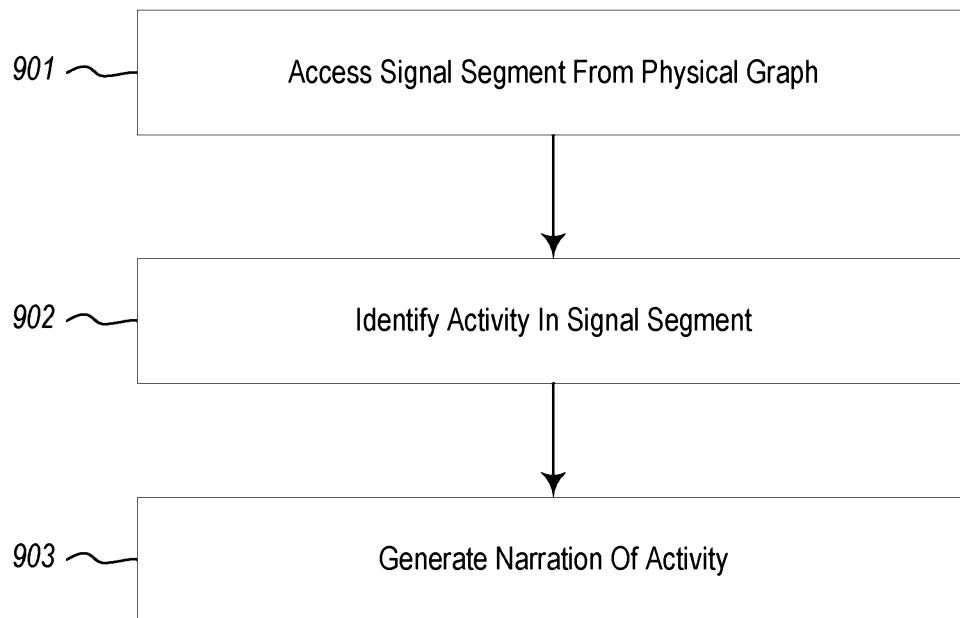
5/8

500**Figure 5**600**Figure 6**

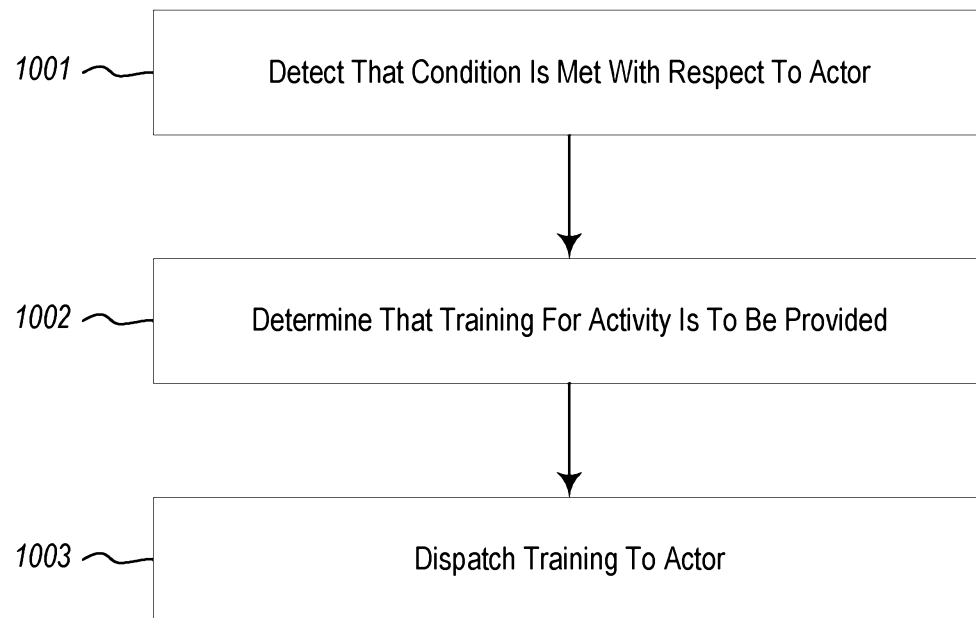
6/8

700**Figure 7**800**Figure 8**

7/8

900**Figure 9**

8/8

1000**Figure 10**