



(51) International Patent Classification:
G06K 9/66 (2006.01)

(21) International Application Number:
PCT/US2019/051874

(22) International Filing Date:
19 September 2019 (19.09.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/734,491 21 September 2018 (21.09.2018) US

(71) Applicant: POSITION IMAGING, INC. [US/US]; 155 Fleet Street, Portsmouth, New Hampshire 03801 (US).

(72) Inventors: CHAKRAVARTY, Narasimhachary Nalana; 934 Portland Ave, Rollingsford, New Hampshire 03869 (US). MIN, Guohua; 8 Millstream Drive, Exeter, New Hampshire 03833 (US). HILL, Edward L.; 305 W. Main Street, Conway, New Hampshire 03818 (US). BILBREY, Brett; 1141 Robin Way, Sunnyvale, California 94087 (US).

(74) Agent: ZALE, Lawrence P.; Schmeiser, Olsen & Watts LLP, 33 Boston Post Road West, Suite 410, Marlborough, Massachusetts 01752 (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).

(54) Title: MACHINE-LEARNING-ASSISTED SELF-IMPROVING OBJECT-IDENTIFICATION SYSTEM AND METHOD

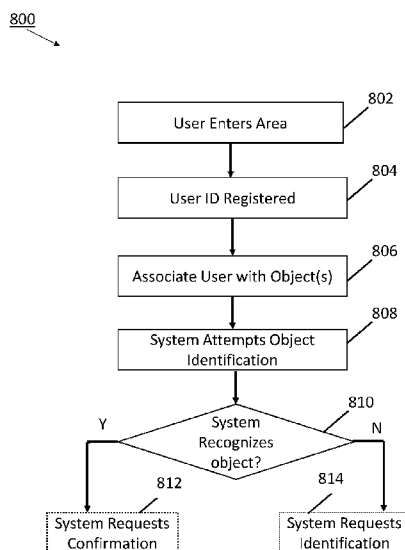


FIG. 8

(57) Abstract: A system and method of identifying and tracking objects comprises registering an identity of a person who visits an area designated for holding objects, capturing an image of the area designated for holding objects, submitting a version of the image to a deep neural network trained to detect and recognize objects in images like those objects held in the designated area, detecting an object in the version of the image, associating the registered identity of the person with the detected object, retraining the deep neural network using the version of the image if the deep neural network is unable to recognize the detected object, and tracking a location of the detected object while the detected object is in the area designated for holding objects.



Published:

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

MACHINE-LEARNING-ASSISTED SELF-IMPROVING OBJECT-IDENTIFICATION SYSTEM AND METHOD

CROSS-REFERENCE TO RELATED APPLICATION

This application claims priority to and the benefit of co-pending U.S. Provisional Application No. 62/734,491, titled “Machine-Learning-Assisted Self-Improving Object-identification System and Method”, filed on September 21, 2018, the entirety of which provisional application is incorporated by reference herein for all purposes.

FIELD OF THE INVENTION

The invention relates generally to machine-learning systems and methods for identifying objects.

BACKGROUND

Various consumer goods and retail operations are attempting to improve customers’ shopping experience by automating the purchase and checkout process. Such automation entails deploying systems that can identify what items a customer has taken from a shelf. Some systems employ video monitoring and image processing techniques to identify those items. However, the proper detection and identification of an item in captured images can be affected by various factors, for example, lighting conditions, shadows, obstructed views, and the location and position of the item on the shelf. Inconsistent results render such systems ineffectual.

SUMMARY

All examples and features mentioned below can be combined in any technically possible way.

In one aspect, the invention is related to an object-identification system comprising an image sensor configured to capture images of objects disposed in an area designated for holding objects, a deep neural network trained to detect and recognize objects in images like those objects held in the designated area, and a controller in communication with the image sensor to receive images captured by the image sensor and with the deep neural network. The controller includes one or more processors configured to register an identity of a person who visits the area designated for holding objects, to submit an image to the deep neural network, to associate the registered identity of the person with an object detected in the image submitted to the deep neural network, to retrain the deep neural network using the submitted image if the deep neural network is unable to recognize the object detected in the submitted image, and to track a location of the detected object while the detected object is in the area designated for holding objects.

The controller may be further configured to acquire labeling information for the detected object in response to the deep neural network being unable to recognize the detected object in the submitted image, to associate the labeling information with the version of the image submitted to the deep neural network, and to store the version of the image and associated labeling information in an image database used to retrain the deep neural network. A human-input acquisition module may be configured to acquire the labeling information from a user in response to a request from the controller

when the deep neural network is unable to recognize the detected object in the image in the submitted image.

The controller may be further configured to find an area within the image in which a change appears if the deep neural network is unable to recognize the detected object, to produce a version of the image that focuses upon the area of change, and to submit the version of the image to the deep neural network to determine whether the deep neural network is able to recognize the detected object in the second version of the image. The controller may be further configured to acquire labeling information for the detected object irrespective of whether the deep neural network recognizes the detected object in the submitted version of the image, to associate the acquired labeling information with the version of the image submitted to the deep neural network, and to store the version of the image and associated labeling information in an image database used to retrain the deep neural network. In addition, the controller may be further configured to acquire the labeling information from the deep neural network when the deep neural network recognizes the detected object in the submitted version of the image.

The deep neural network may be a first deep neural network, and the system may further comprise a second deep neural network configured to operate in parallel to the first deep neural network. Each of the first and second deep neural networks produce an output based on image data obtained from the image, wherein the image data obtained by the first deep neural network are different from the image data obtained by the second deep neural network.

The object-identification system may further comprise a depth sensor with a field of view that substantially matches a field of view of the image sensor. The depth sensor acquires depth pixel value of images within its field of view, wherein a depth pixel value and less than three pixel values taken from the group consisting of R (red), G (green), and B (blue) are submitted as image data to the deep neural network when the image is submitted to the deep neural network during training or object recognition.

The deep neural network may reside on a remote server system, and the controller may further comprise a network interface to communicate with the deep neural network on the server system.

In another aspect, the invention is related to a method of identifying and tracking objects. The method comprises the steps of registering an identity of a person who visits an area designated for holding objects, capturing an image of the area designated for holding objects, submitting a version of the image to a deep neural network trained to detect and recognize objects in images like those objects held in the designated area, detecting an object in the version of the image, associating the registered identity of the person with the detected object, retraining the deep neural network using the version of the image if the deep neural network is unable to recognize the detected object, and tracking a location of the detected object while the detected object is in the area designated for holding objects.

The method may further comprise acquiring labeling information for the object detected in the version of the image in response to the deep neural network being unable to recognize the detected

object in the version of the image, associating the labeling information with the version of the image, and storing the version of the captured image and associated labeling information in an image database used to retrain the deep neural network. The step of acquiring labeling information for the object detected in the version of the image may be in response to the deep neural network being
5 unable to recognize the detected object in the version of the image comprises acquiring the labeling information from user-supplied input.

The method may further comprise finding an area within the version of the image in which a change appears when the deep neural network is unable to recognize the object detected in the first version of the image, producing a second version of the image that focuses upon the found area of
10 change, and submitting the second version of the image to the deep neural network to determine whether the deep neural network can recognize the detected object in the second version of the image. The method may further comprise acquiring labeling information for the object detected in the first version of the image irrespective of whether the deep neural network recognizes the detected object in the second version of the image, associating the labeling information with the first version of the
15 image, and storing the first version of the captured image and associated labeling information in an image database used to retrain the deep neural network. The step of acquiring labeling information for the object detected in the version of the image may comprise acquiring the labeling information from the deep neural network when the deep neural network recognizes the detected object in the version of the image.

20 The step of submitting a version of the image to the deep neural network may comprises submitting a depth pixel value and less than three pixel values taken from the group consisting of R (red), G (green), and B (blue) as image data to the deep neural network.

The method may further comprise the step of submitting image data, acquired from the version of the image, to a first deep neural network and a second deep neural network in parallel, wherein the
25 image data submitted to the first deep neural network are different from the image data submitted to the second deep neural network.

In another aspect, the invention is related to a sensor module comprising an image sensor configured to capture an image within its field of view and a depth sensor having a field of view that substantially matches the field of view of the image sensor. The depth sensor is configured to acquire
30 estimated depth values for an image captured by the depth sensor. The sensor module further comprises a controller in communication with the image sensor and depth sensor to receive image data associated with the image captured by the image sensor and estimated depth values associated with the image captured by the depth sensor. The controller includes one or more processors configured to register an identity of a person who visits an area designated for holding objects, to
35 submit the image data associated with the image captured by the image sensor and the estimated depth values associated with the image captured by the depth sensor to a deep neural network trained to detect and recognize objects in images like those objects held in the designated area, to associate the

registered identity of the person with an object detected in the image data and estimated depth values submitted to the deep neural network, and to save a version of the images captured by the image sensor and the depth sensor for use in subsequent retraining of the deep neural network if the deep neural network is unable to recognize the detected object.

5 The controller may further comprise a cloud interface to communicate with the deep neural network over a network.

The sensor module may further comprise a human-input acquisition module configured to acquire the labeling information from a user in response to a request from the controller when the deep neural network is unable to recognize the detected object based on the submitted image data and estimated
10 depth values.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of this invention may be better understood by referring to the following description in conjunction with the accompanying drawings, in which like numerals indicate like structural elements and features in various figures. The drawings are not necessarily to
15 scale, emphasis instead being placed upon illustrating the principles of the invention.

FIG. 1 shows an embodiment of a machine-learning-assisted object-identification system including a controller and an image sensor having a field of view of a support surface that is designated for holding objects.

FIG. 2 is a block diagram of an embodiment of the controller of FIG. 1

20 FIG. 3 is a functional block diagram of an embodiment of the machine-learning-assisted object-identification system including an artificial intelligence module in communication with a local machine-learning module.

FIG. 4 is a functional block diagram of an embodiment of the machine-learning-assisted object-identification system including an artificial intelligence module in communication with a remote
25 machine-learning module.

FIG. 5 is a functional block diagram of another embodiment of the machine-learning-assisted object-identification system including an artificial intelligence module in communication with a remote machine-learning module.

30 FIG. 6 is a block diagram of an embodiment of an object-tracking module including a deep neural network or DNN.

FIG. 7 is a block diagram of an embodiment of the computer-vision module.

FIG. 8 is a flow chart of an embodiment of a process for machine-learning-assisted object identification.

35 FIG. 9 is a flow chart of an embodiment of a process for identifying an object in a captured image and an optional multi-pass authentication process.

FIG. 10 is a flow chart of an embodiment of a process for retraining the DNN with images in which the DNN does not initially identify an object.

DETAILED DESCRIPTION

Object-identification systems and methods described herein combine computer vision, machine learning, and a deep neural network (DNN) to enable the accurate identification and tracking of people and objects. Initially, the DNN may be a blank slate and incapable of object identification without human assistance or it can be trained with a predetermined set of images to give it a baseline. To give baseline object identification capabilities to the DNN, a human has to train the DNN with a predetermined training set of images. After its initial training, the DNN's ability to identify objects continuously improves because of subsequent trainings. These subsequent trainings are based on images in which the DNN could not initially identify an object. Objects in these images have become identifiable, and thus valuable for retraining the DNN, because of human-supplied information that identifies objects in the images or because of a multi-authentication process that focuses the DNN's detection efforts on a region in the images where change has been detected.

FIG. 1 shows an embodiment of a machine-learning-assisted object-identification system 100 having at least one sensor module 102. Each sensor module 102 includes a controller 104 that is in communication with one or more color (e.g., RGB) image sensors 106, and, optionally, as shown in phantom, one or more depth sensors 108 and a light source 110, and, optionally, a wireless radio-frequency (RF) transceiver (not shown). The controller 104 implements a deep neural network (DNN) 112 for use in object recognition and a computer-vision module 114 for detecting changes in images, as described later in more detail.

In one embodiment, each sensor module 102 is a self-contained electronic unit capable of registering persons who visit the object-holding area, capturing images, image processing, detecting objects, machine-learning-assisted self-improving object recognition, object tracking, and, when so configured, providing light guidance. In other embodiments, one or more of these functions takes place remotely (i.e., not at the sensor module); for example, the functions of object detection, machine-learning-assisted self-improving object recognition, and object tracking can occur at a remote computing site with which the sensor module is in communication over a network.

The sensor module 102 may be deployed in a fixed position near a support surface 116 in an object-holding area, or it may be mobile, embodied in a mobile device. As an example of a fixed deployment, the sensor modules 102 may drop down from the ceilings in a surveillance configuration so that all corners of an enterprise site are covered. These sensor modules are small and non-intrusive and can track the identifications and paths of individuals through the enterprise, for example, as described in U.S. Pat. Pub. No. US-2018-0164112-A1, published June 14, 2018, the entirety of which application is incorporated by reference herein.

Mobile embodiments of the sensor module include, but are not limited to, a smartphone, tablet computer, wearable computing device, or any other portable computing device configured with one or more processors, an RGB camera, wireless communication capabilities, an optional depth sensor, an optional light source, and software for performing the image processing, object detecting, tracking,

and recognizing, self-improving machine learning, and optional light guidance functions described herein. The software can be embodied in a downloaded application (app) that can be stored on the mobile device. Being portable, a person or machine can, in effect, carry an object-identification device capable of recognizing objects captured by the camera(s) of the mobile device. For example, a
5 person with such a device can run the software, approach a table (i.e., support surface) holding various objects, point the device (i.e., its camera(s)) at each object, capture an image of an object, and be told the type (identity) of the object. To obtain the identity of the object, the mobile device may communicate with a remote server that hosts the DNN, sending the image to the remote server, and receiving the identity of the object.

10 Each image sensor 106, which may also be referred to herein as an optical sensor, provides color information; each depth sensor 108 provides estimated depth for each pixel of a captured image. The image sensor 106 and depth sensor 108 may be embodied in a single camera, such as, for example, Microsoft's Kinect™, or be embodied in separate cameras. The image and optional depth sensors are disposed to face the support surface 116. Examples of the support surface include, but are not limited to, desktops, tables, shelves, and floor space. In general, the support surface is disposed in or at an
15 object-holding area. The object-holding area can be, for example, a supermarket, warehouse, inventory, room, closet, hallway, cupboards, lockers, each with or without secured access. Examples of identified and tracked objects include, but are not limited to, packages, parcels, boxes, equipment, tools, food products, bottles, jars, and cans. (People may also be identified and tracked.) Each image
20 sensor 106 has a field of view (FOV) that covers a portion of, or all the area occupied by the support surface 116; the field of view of an optional depth sensor matches at least that of an image sensor. Each separate sensor has its own perspective of the area and of the objects placed on the support surface 116.

25 The controller 104 may be configured to control the light source 110 to provide light guidance to objects located on the support surface 116 or to certain regions of the support surface, depending upon the object or region of interest. Examples of the light source 110 include, but are not limited to, lasers, projectors, LEDs, light bulbs, flashlights, and lights. The light source 110 may be disposed on or remote from and directed at the support surface 116.

30 A display 118 may be included in the object-identification system 100, to provide, for example, a visual layout of the objects on the support surface, visual guidance to objects or regions on the surface, and a user interface for use by persons who enter and leave the object-holding area. The display 118 may be conveniently located at the threshold of or within the holding area. The display 118 may be part of an electronic device (e.g., a computer, smartphone, mobile device) configured with input/output devices, for example, a physical or virtual keyboard, keypad, barcode scanner,
35 microphone, camera, and may be used to register the identities of persons entering the object-holding area and/or to scan object labels.

The controller 104 may also be in communication with one or more servers 120 (i.e., server system) over a network connection. These server(s) 120 may perform third-party services, such as “cloud services,” or be implemented locally or onsite at the enterprise. As used herein, the “cloud” refers to software and services that run on a remote network, such as the Internet, instead of at the sensor module or at a local computer. The cloud may be public, private, or a combination thereof. An example of cloud services suitable for the principles described herein is Azure™ cloud services provided by Microsoft® of Redmond, WA. The server(s) 120 can run a virtual machine that provides the cloud services required by the sensor module 102.

During operation of the object-identification system 100, persons arrive at the object-holding area to perform any one or more of at least four object handling activities, including depositing an object, removing an object, moving an object to another spot in the holding area, or alerting personnel of an object warranting inspection. In general, the object-identification system registers the identities of persons who arrive at the holding area (i.e., who interact with the object-identification system) and associates each registered person with one or more objects that the person is handling. Using image processing techniques, the object-identification system continuously monitors and acquires real-time image data of the holding area. From the real-time image data, the object-identification system detects when each such object is placed on the support surface 116, moved to another region of the support surface, or removed from the support surface. Techniques for detecting and tracking objects disposed on a support surface in a holding area can be found in U.S. Pat. Appln. No. 15/091,180, filed April 5, 2016, titled “Package Tracking Systems and Methods,” the entirety of which patent application is incorporated by reference herein. In addition, the object-identification system may identify a perishable item and send a notification to staff of its expiration. Or the object-identification system can recognize damaged goods on a shelf and notify staff accordingly. In response to the notifications, staff can then inspect the item in question to remove if past its expiration date or confirm the extent of damaged packaging.

The object-identification system further recognizes each object on the support surface or involved in a handling activity. Object recognition serves to identify the type of object detected and tracked (e.g., a package from a certain carrier, a jar of pickles, a microscope). Such object recognition may involve human interaction to initially identify or to confirm, correct, or fine tune the recognition of a given object. The object-identification system employs machine-learning techniques to improve its object recognition capabilities. Recognition of a given object can facilitate the tracking of the object while the object is in the holding area, serving to confirm the presence or movement of the object.

Upon occasion, the sensor module 102 will capture an image for which object recognition falls below a threshold, namely, the object-identification system is unable to recognize an object in the image. Despite being unable to recognize the object (at least initially), the object-identification system can still track the object, namely, its initial placement and any subsequent location within the holding area, based on visual characteristics of the object. The unidentifiable image is retained for

purposes of later retraining of the DNN 112 so that the DNN will become able to recognize a previously unrecognizable object when that object is present in subsequently processed images.

Human interaction with the object-identification system, through voice recognition, gesture recognition, or keyboard input, can specifically identify an object in an unidentifiable image, giving the image a proper label. An example of gesture recognition is a person holding up three fingers to identify the object as type number 3, where the object-identification system has stored the association of a three-finger gesture with a specific object (e.g., three fingers correspond to a microscope). After an object in the previously unidentifiable image becomes recognized, with the help of the human input, the image and associated proper label are stored in an image database 122. The object-identification system 100 uses these stored images and labels to retrain the deep neural network 112. By retraining the deep neural network with previously unidentifiable images, now made identifiable by human-provided information, the neural network 112 increasingly grows “smarter”. Over time, the probability of the neural network recognizing objects in later captured images approaches one hundred percent.

The image database 122 may be kept in local storage 124, accessed through a central computer 126 in proximity of the sensor module 102. In this embodiment, the central computer 126 provides access to the image database 122 for all deployed sensor modules 102. In another embodiment, shown in phantom in FIG. 1, the image database 122 is stored in remote storage 128, for example, in “the cloud”, with which each sensor module 102 is in communication through the server(s) 120. In addition to those initially unidentifiable images, the image database 122 also holds the initial training set of images.

FIG. 2 shows an embodiment of the controller 104 of FIG. 1. The controller 104 includes one or more processors 200, examples of which include, but are not limited to, image processors, central processing units, graphics processing units, each of standard or custom design. The one or more processors 200 are in communication with memory 202. In the instance of multiple processors 200, such processors may be located at different sites (e.g., one processor disposed locally (i.e., at the sensor module) and another disposed remotely (e.g., in “the cloud”). Similarly, the memory 202 can be disposed locally, remotely, or a combination thereof.

The one or more processors 200 are in communication with a video interface 204, an optional light source interface 206, an optional audio interface 208, a network interface 210, and interfaces 212 to I/O components (e.g., the display 118). By the video interface 204, the controller 104 communicates with each image sensor 106 and depth sensor 108, if any, in the sensor module 102; by the light source interface 206, the controller 104 controls activation of the light source 110, and, depending upon the type of light source, the direction in which to point an emitted light beam; by the audio interface 208, the controller 104 communicates with audio devices that capture or play sound.

In addition to conventional software, such as an operating system and input/output routines, the memory 202 stores program code for configuring the one or more processors 200 to implement the

deep neural network (DNN) 112, and to perform personnel registration 214, object detection 216 in images, object tracking 218 in the holding area, object recognition 220 in images, neural network training 222, image-preprocessing 224, change tracking 226 in images, and, optionally, light guidance 228. The one or more processors 200 and memory 202 can be implemented together or individually, on a single or multiple integrated circuit (IC) devices. In addition, the program code stored in memory 202 can reside at different sites. For example, the program code for implementing the DNN 112 can reside at a remote location (e.g., on the cloud) while the program code for user recognition can reside and execute locally (i.e., on the sensor module).

In brief overview, the program code for personnel registration 214 records the identities and activities of individuals who use the object-identification system 100 and associates such individuals with the objects they affect; the program code for object detection 216 uses image-processing techniques to detect the presence of objects in images; the program code for object tracking 218 tracks the locations of detected objects within the holding area, the program code for object recognition 220 employs the DNN 112 to recognize (i.e., identify or classify) objects in images; the program code for neural network training 222 trains the DNN 112 to become capable of recognizing particular types of objects; the program code for image pre-processing 224 applies image editing techniques to captured images to improve object detection and recognition efforts in such images; the program code for change tracking 226 detects changes in images and assists in labeling images; and, optionally, the program code for light guidance 228 guides humans to objects and/or locations in the object-holding area using the light source 110. As later described in more detail, various elements or functionality of the controller 104 may reside remotely; that is, in some embodiments, some elements or functionality of the controller 104 are not part of the sensor module 102 (FIG. 1), but reside remotely (e.g., in “the cloud”).

FIG. 3 shows a functional block diagram of an embodiment of the object-identification system 100 including an artificial-intelligence (AI) module 300 in communication with a machine-learning module 302. The AI module 300 includes an image-acquisition module 304, an image-preprocessing module 306, an object-tracking module 308, a human-input-acquisition module 310, and tracking quality measurement module (QMM) 312. The machine-learning module 302 includes the local storage 124 (which maintains the image database 122, described in connection with FIG. 1) and a deep neural network (DNN) trainer 314.

The image acquisition-module 304 of the AI module 300 is configured to acquire images from the image sensor 106 and optional depth sensor 108. Captured images pass to the image-preprocessing module 306, and the image-preprocessing module 306 forwards the images to the object-tracking module 308. The image-preprocessing module 306 sends each image (line 316) to the computer-vision module 114 and copy of that image (line 318) to the DNN 112 (alternatively, the computer-vision module 114 receives the copy of the image).

In general, the object-tracking module 308 is configured to detect objects in images, to track such objects, and to perform object recognition using the DNN 112 of FIG. 1. The results (line 320) produced by the object-tracking module 308 pass to the tracking QMM 312, which uses thresholds (or other criteria) to establish whether an object has been recognized in an image. The results comprise
5 two types of information: 1) the area in the image that the DNN deems to have found an object (or areas for multiple objects, depending on the image); and 2) a list of probabilities of each type of object considered to be in that area or areas.

Based on the information received from the object-tracking module 308, the QMM 312 determines whether the DNN 112 was successful identifying an object (or objects) in an image. If
10 successful, the QMM 312 signals (line 322) success. The controller 104 can receive this success signal and respond to the signal accordingly, depending upon the end-user application that seeks to determine the identification of objects, such as a package-tracking application.

If an object is not identifiable within an image, the QMM 312 notifies (line 324) the computer-vision module 114. The computer-vision module 114 optionally sends an image (line 326) to the
15 DNN 112; this image is derived from the original image and focuses on a region in the original image in which a change was detected. The DNN 112 may attempt to identify an object in this focused image (line 326), that is, the DNN 112 performs a second pass. If the DNN is unsuccessful during the second pass, the QMM 312 sends a request (line 327) to the human-input-acquisition module 310, seeking labeling information for the unidentifiable object in the original image. Irrespective of the
20 success or failure of the DNN 112 to recognize an object in this focused image, the computer-vision module 114 sends (line 328) the original image within which an object was not initially recognized to the local storage 124. The image being stored is joined/associated (box 330) with a human-provided label (line 332) from the human-input-acquisition module 310 or with a label (line 334) produced by the DNN 112 (line 320), sent to the QMM 312, and then forwarded by the QMM 312. The DNN
25 trainer 314 uses those images in the local storage 124 and their associated ID information (i.e., labels) to retrain (line 336) the DNN 112.

Each sensor module 102 (FIG. 1) can be configured to locally provide the functionality of the AI module 300, and the central computer 126 (FIG. 1) is configured to provide the functionality of the machine-learning module 302. Each sensor module 102 (FIG. 1) is in communication with the central
30 computer 126, which provides the DNN training based on the images in the image database 122. Accordingly, different sensor modules will develop equivalent image recognition capabilities.

FIG. 4 shows a functional block diagram of another embodiment of the object-identification system 100 having an artificial intelligence (AI) module 400 in communication with a machine-learning module 402. In general, the object-identification system of FIG 4 differs from that of FIG. 3
35 in that the machine-learning module 402 resides remotely, that is, on a network that is in the cloud or on the Internet. In addition, in this embodiment, the remote machine-learning module 402 includes

the DNN 112. Accordingly, the DNN computations occur on a remote server (e.g., server 120 of FIG. 1), performed, for example, by third-party cloud services.

Specifically, the AI module 400 includes the image-acquisition module 304 (FIG. 3), an image-preprocessing module 306 (FIG. 3), an object-tracking module 404, the human-input-acquisition module 310 (FIG. 3), and the tracking QMM 312. The object-tracking module 404 includes the computer-vision module 114 and a cloud-interface module 406. The remote machine-learning module 402 maintains the image database 122 (FIG. 1) in cloud (i.e., remote) storage 128 and includes the DNN 112 and the DNN trainer 314 used to train and retrain (line 336) the DNN 112.

The computer-vision module 114, image-acquisition module 304, image-preprocessing module 306, object-tracking module 404, human-input-acquisition module 310, tracking QMM 312, cloud (i.e., remote) storage 128, DNN trainer 314, and DNN 112 operate like their counterpart modules in FIG. 3. A difference is that the AI module 400 uses its cloud-interface module 406 to transfer images, received from the image-acquisition module 304, over the network to the remote machine-learning module 402, to be used by the DNN 112 in its object detecting, tracking, and recognition algorithms.

To the image-acquisition module 304 and the computer-vision module 114, the cloud-interface module 406 is effectively a “virtual DNN”, that receives input from these two modules 304, 114 as though it were the DNN, and forwards such input (line 408) to the DNN 112 residing remotely. Accordingly, processor(s) on the remote server 120 perform the DNN computations. In addition, the machine-learning module 402 is configured to return the results (line 320) produced by the DNN 112 to the cloud-interface module 406 of the AI module 400. The cloud-interface module 406 forwards these results (line 410) to the QMM 312, which determines from the results whether the DNN 112 has recognized an object in the image.

If the QMM 312 determines the DNN 112 was successful identifying an object (or objects) in an image, the QMM 312 signals (line 322) success. If an object is not identifiable within an image, the QMM 312 notifies (line 324) the computer-vision module 114. The computer-vision module 114 optionally sends an image (line 326) to the cloud interface 414, for transmission to the remote DNN 112. This image is derived from the original image and is focused on a region in the original image in which the computer-vision module 114 detected change. The DNN 112 may attempt to identify an object in this focused image. If the DNN attempts but is unsuccessful during the second pass, the QMM 312 sends a request (not shown) to the human-input-acquisition module 310, seeking labeling information for the unidentifiable object in the original image.

Irrespective of the success or failure of the DNN 112 to recognize an object in this focused image during the DNN’s second attempt, the computer-vision module 114 forwards (line 328) the original image (or an edited version of the original image) to the cloud storage 128. The image to be stored is joined or associated with (box 330) a human-provided label (line 332) acquired by the human-input-acquisition module 310 (in the event of a DNN failure) or with a label (line 320) produced by the DNN 112 on a successful second pass and forwarded by the QMM 312 (in the event of the DNN

success). The DNN trainer 314 uses those images in the remote storage 128 and their associated ID information (i.e., labeling information) to retrain the DNN 112.

Each sensor module 102 (FIG. 1) can be configured to locally provide the functionality of the AI module 400, and the server system 120 (FIG. 1) is configured to provide the functionality of the machine-learning module 402. In this embodiment, all sensor modules 102 that are in communication with the remote machine-learning module 402 share the DNN 112 and image database 122 and contribute images for future re-training of the DNN 112. In general, by having access to the same image database 122 and DNN 112, all sensor modules will develop equivalent image recognition capabilities.

FIG. 5 shows a functional block diagram of another embodiment of the object-identification system 100 having an artificial intelligence (AI) module 500 in communication with a machine-learning module 502. Each sensor module 102 (FIG. 1) can be configured to locally provide the functionality of the AI module 500, and the server system 120 (FIG. 1) is configured to provide the functionality of the machine-learning module 502. In general, the object-identification system of FIG. 5 differs from that of FIG. 4 in that certain functionality of the AI module 400 in FIG. 4 occurs remotely. Specifically, the image-preprocessing module 306, the object-tracking module 308, and the QMM 312 are part of the remote machine-learning module 502, which also includes the DNN 112, remote storage 128, and DNN trainer 314. The AI module 500 includes a cloud interface 504, the image-acquisition module 304, and the human-input-acquisition module 310. The cloud interface 504 is in communication with the image-acquisition module 304 to forward (line 450) input images received therefrom to the remote image-preprocessing module 306. The cloud interface 504 also receives a successful result indicator along with results (line 322) from the remote QMM 312 when the DNN 112 is able to recognize one or more objects in the input image; the cloud interface 504 outputs the result for use by an application of the object-identification system 100.

If an object is not identifiable within an image, the QMM 312 signals (line 324) the computer-vision module 114. In response to the “DNN FAILS” signal, the computer-vision module 114 may send an image (line 326), derived from the original image (or an edited version of it) that is focused on a region in the original image in which the computer-vision module 114 detects a change, to the DNN 112 for an attempt to identify an object in this focused image, in effect, performing a second pass at authentication. The DNN 112 sends (line 320) the results of this second attempt to the QMM 312.

Irrespective of the success or failure of the DNN 112 to recognize an object in the focused image during the second attempt, the remote computer-vision module 114 forwards (line 328) the original image (or an edited version thereof), in which the DNN 112 was initially unable to recognize an object, to the cloud storage 128.

If an object is not identifiable within this focused image, the QMM 312 signals the AI module 500 (line 327), telling the AI module 500 to request human input. When the human-input-acquisition

module 310 receives the human input, the cloud interface 504 sends (line 332) a human-input label to the cloud storage 128. Before being stored, the human input label (line 332) is combined or associated with (box 330) the image coming from the remote computer-vision module 114.

If an object is identifiable within the focused image, the QMM 312 sends a label (line 334) produced by the DNN 112 that is combined or associated with (box 330) the image sent to the cloud storage 128 by the computer-vision module 114. As previously described, the DNN trainer 314 uses those images and their associated labels in image database 122 maintained in the remote storage 128 to retrain (line 336) the DNN 112.

FIG. 6 shows an embodiment of an object-tracking module 600, including the image pre-processing module (or image preprocessor) 306 and the computer-vision module 114, for use in connection with the embodiments of object-identification systems described in FIG. 3 and FIG. 5. The object-tracking module 600 includes the DNN 112. In an embodiment where the DNN 112 is remote with respect to the object-tracking module, such as in FIG. 4, the object-tracking module 600 has a cloud-interface module 406 (FIG. 4) that operates as a virtual DNN and communicates with the actual DNN.

In one embodiment, the DNN 112 has a deep learning architecture, for example, a deep convolutional neural network, having an input layer 602, an output layer 604, and multiple hidden layers (not shown). Hidden layers may comprise one or more convolutional layers, one or more fully connected layers, and one or more max pooling layers. Each convolutional and fully connected layer receives inputs from its preceding layer and applies a transformation to these inputs based on current parameter values for that layer. Example architectures upon which to implement a deep learning neural network include, but are not limited to, the Darknet Open source Deep Neural Net framework available at the website pjreddie.com and the Caffe framework available at the website caffe.berkeleyvision.org.

The DNN 112 is involved in two processes: object detection/recognition and training. For purposes of object detection and recognition, images 606 are provided as input 608 to the DNN 112 from the image-acquisition module. The images 606 include color images (e.g., RGB) and, optionally, depth images. Color and depth images captured at a given instant in real-time are linked as a pair. Such images may pass through the image preprocessor 306, which produces image data 608 based on the processed images. The image preprocessor 306 may or may not modify an image before the image 606 passes to the DNN. In one embodiment, the image preprocessor 306 is configured to apply one or more image-editing techniques determined to enhance the DNN's ability to detect objects in images by making such images robust (i.e., invariant) to illumination changes. For RGB, one pre-processing algorithm uses a series of steps to counter the effects of illumination variation, local shadowing, and highlights. Steps in the algorithm include gamma correction, difference of Gaussian filtering, masking and contrast equalization. Depth data can be noisy and can have missing data depending on the circumstances under which depth data are captured. Ambient light and highly

reflective surfaces are major factors of noise and missing data. This pre-filtering ensures that these artifacts are corrected for and that the data are well preserved. Pre-processing steps include ambient light filtering, edge-preserving smoothing, Gaussian blurring, and time-variant blurring. When depth images and RGB images both pass to the image preprocessor 306, the image preprocessor performs a
5 blending transformation process that blends the RGB data with the depth data to produce image data 608. Examples of blending transformation processes include, but are not limited to, blending by concatenation or blending by interleaving, both of which are described in more detail below.

The image data 608 passes to the input layer 602 of the DNN 112. (Though not shown in FIG. 6, any label information obtained by human interaction also passes to the input layer 602 for training,
10 object recognition, or both). Both the image preprocessor 306 and the computer-vision module 114 are configured to provide image data 608 to the DNN 112. The image data 608 includes pixel values for each of RGB (Red, Green, Blue) acquired from the color images and, optionally, for D (Depth) taken from depth images. FIG. 6 illustrates blending by concatenation, wherein each color and pixel value has its own channel, and the RGB image is blended with the depth image by concatenating the
15 channels. In those embodiments where D values are submitted, less than all color values R, G, or B, may be submitted. For example, R+D, G+D, B+D, R+G+D, R+B+D, G+B+D are instances of where less than all three color (RGB) values are submitted as input together with the D value, each submitted value having its own channel.

Alternatively, blending by interleaving can blend the RGB image with the depth image. In this
20 blending technique, instead of concatenating the RGB and depth images and gaining channels, the channels of both images are blended in a manner that retains the original structure, that is, the number of channels in the resulting image do not increase after blending from the number of channels in the original RGB image. One such example follows:

Consider an eight-bit three-channel RGB image, that is, the R-channel has eight bits, the G-
25 channel has eight bits, and the B-channel has eight bits. Further, consider that the depth image is a single channel of 16-bit data; that is the D-channel has 16 bits.

One method of combining data from multiple dimensions (i.e., channels) and packing the data into fewer dimensions (i.e., channels) is the Morton Number Interleaving.

For example, a color pixel value [R, G, B] of [255, 125, 0] has an eight-bit binary representation
30 of [11111111, 01111101, 00000000], where the three eight-bit values represent the three eight-bit R, G, and B channels, respectively.

For the 16-bit depth value, three eight-bit values are derived. The first eight-bit value, referred to as D1, entails a conversion of the 16-bit value to an eight-bit value. This conversion is done by
normalizing the decimal equivalent of the 16-bit depth value and multiplying the normalized value by
35 the maximum value of an eight-bit number (i.e., 255). For example, consider an original 16-bit depth value [D] that has a decimal value of [1465]. Normalizing the decimal value [1465] entails dividing this decimal value by the maximum decimal value that can be represented by 16 bits, namely [65025].

Accordingly, the multiplied, normalized decimal value for $D1 = (1465/65025) * 255 = 6$ (rounded up). The eight-bit binary representation of $D1$ is [00000110].

The next two bytes are obtained by partitioning the original 16-bit depth value [D] into two eight-bit bytes, called $D2$ and $D3$. For example, the previously noted 16-bit depth value [D] of [1465] has a binary representation of [0000010110111001]. The 8-bit $D2$ byte corresponds to the first byte of the 16-bit depth value [D], which is [00000101], and the 8-bit $D3$ byte corresponds to the second byte of the 16-bit depth value [D], which is [10111001]. Accordingly, $[D2, D3] = [00000101, 10111001]$.

The three bytes $[D1, D2, D3]$ derived from the original depth value [D] are [00000110, 00000101, 10111001]. As previously mentioned, the three-channel, eight-bit RGB values are [11111111, 01111101, 00000000].

Morton order interleaving produces a 16-bit, three-channel image from the three channels of depth values bytes $[D1, D2, D3]$ and the three channels of RGB values $[R, G, B]$ bytes by appending the depth values to the RGB values as such: $[RD1, GD2, BD3]$. With regards to the previous example, the Morton order interleaving produces three 16-bit channels of [1111111100000110, 0111110100000101, 0000000010111001]. The technique executes for each pixel of the corresponding images 606 (i.e., RGB image and its associated depth image). The result is a three-channel image that has both depth and color information. It is to be understood that Morton order interleaving is just an example of a technique for interleaving depth data with color data for a given pixel; other interleaving techniques may be employed without departing from the principles described herein.

As with the blending by concatenation technique, less than all color values $R, G,$ or $B,$ may be interleaved with a depth value. For example, $R+D, G+D, B+D, R+G+D, R+B+D, G+B+D$ are instances of where less than all three color (RGB) values are submitted as input together with a D value. In these cases, there is a separate channel for each interleave of color and depth. When less than three RGB channels are used, any of the $D1, D2, D3$ depth channels can serve for interleaving. For example, the combinations such as $R+D, G+D,$ and $B+D$ each require only one channel; combinations such as $R+G+D, R+B+D, G+B+D$ each have two channels. If only one RGB channel is used, $D1$ is the preferred choice, because the $D1$ depth channel contains the whole depth information. If two color channels are used, then two depth channels are used in the interleaving: for example, $D2$ and $D3$ ($D2$ and $D3$ together have the whole depth information). To illustrate, again using the color pixel value $[R, G, B]$ of [255, 125, 0] and the original depth value of [1465], the combination of $R+G+D$ produces the following 16-bit two channel $[RD2, GD3]$ input data: [1111111100000110, 0111110100000101], where $D2$ and $D3$ are the chosen depth channels. In general, the ability to achieve object detection benefits from having more information available rather than less; accordingly, blending by concatenation, which retains all of the available color and potentially depth data, may produce better detection outcomes than blending by interleaving, which reduces the number

of channels and may use less than all color and depth. Where blending by interleaving may be more advantageous over blending by concatenation is when it comes to training speed.

The output layer 604 produces an output 320 that passes to the QMM, which may be by way of a cloud interface 406 (FIG. 4). The output 320 can include a label for an object detected in the image-under-analysis, an indication of where the object is located in the image, and a value representing the probability (i.e., level of confidence) of an accurate recognition.

The DNN 112 is also in communication with the DNN trainer for purposes of receiving parameter value updates used in retraining.

In one embodiment, the DNN 112 is comprised of two deep neural networks (not shown) that operate in parallel. One neural network receives the R, G, and B pixel values, while the other receives the R, G, B, and D values. Each neural network attempts to recognize one or more objects in the supplied image based on the image data 608 submitted. Each produces an output. The two outputs can be compared and/or combined, for purposes of confirming and/or augmenting each other's determination. For example, consider that the RGB neural network produces a result of having detected one package in a specific area of the image and the RGBD neural network produces a result having detected two packages in the same specific area. A comparison of the probabilities of the two neural networks (and a logic circuit) would declare resolve the difference and finalize the result as either one package or two.

The computer-vision module 114 is in communication with the QMM to receive a "DNN FAILS" signal in the event the DNN 112 fails to successfully recognize an object in the image. Upon receiving such a signal, the computer-vision module 114 outputs (line 328) an image corresponding to the original image in which the DNN could not identify an object. This image can become associated with labeling information 332 supplied by a human (e.g., in response to a prompt from the AI module when the DNN's object identification fails). This combination 610 of labeling information and image passes to storage, where it becomes part of the image database 122. Alternatively, the combination 610 includes the image and labeling information coming (line 334) from the QMM 312 (produced by the DNN 112) when the DNN successfully identifies an object during the second pass.

FIG. 7 shows an embodiment of the computer-vision module 114. In general, the image-acquisition module can capture an image that has a higher resolution than what the DNN 112 needs as an input. For example, the image-acquisition module can acquire an image 606-1 that is 3840 x 2160 pixels, whereas the DNN 112 requires an image of 416 x 416 pixels (e.g., when based on a required calculation speed). Accordingly, the image-preprocessing module 306 down-samples the original image 606-1 to produce a resized image 606-2 that matches the input resolution of the DNN 112. The DNN 112 attempts to detect an object or objects in the resized image 606-2 and sends the results to the QMM 312. The QMM determines from the results whether the DNN 112 has successfully detected one or more objects in the image.

In addition, the image-preprocessing module 306 sends the original image 606-1 and the resized image 606-2 to the computer-vision module 114. The computer-vision module 114 includes a change-tracking module 700 in communication with a change-localization module 702. In one embodiment, the computer-vision module 114 performs a multi-pass authentication process when the DNN 112 fails to detect an object in the image 606-2. In the event of an unsuccessful object detection, the QMM signals the change-tracking module 700, which, in response, executes the change-tracking program code 226 (FIG. 2) to identify an area 704 in the original image 606-1 in which a change appears (with respect to an earlier captured image).

The change-localization module 702 uses this information to produce an image 606-3 that focuses on the region 704 in the original image with the detected change. The focused image 606-3 has a resolution that matches the input resolution of the DNN 112. In order to attain this resolution, the change-localization module 702 may have to reduce or enlarge the size of the region 704 of change. The focused image 606-3 passes to the DNN 112, which attempts to detect an object in this image. The computer-vision module 114 sends the resized image 606-2 to the storage (local or remote) and marks the boundaries of the focus region 704 as those boundaries translate to the resized image 606-2. The boundary information includes a row, column, height, and width of the pertinent region within the resized image 606-2.

Within the storage, the resized image 606-2 is associated with the label name provided by human input (when the DNN fails to recognize an object in the focused image 606-3) or with the label produced by the DNN 112 (when the DNN successfully recognizes an object in the focused image 606-3). The resized image 606-2, the marked boundaries, and label information are used together in subsequent retraining of the DNN 112.

FIG. 8 shows an embodiment of a process 800 for object tracking in accordance with the principles described herein. A person enters (step 802) an area designated for holding objects. Such an area can be, for example, a room for storing packages or a shelf holding boxes of cereal or jars of mayonnaise.

The object-identification system 100 registers (step 804) an identification of the person. The registration can occur automatically, that is, without the person's conscious involvement. For example, a sensor module 102 can wirelessly communicate with a device carried by the person, such as, for example, a key fob or a smartphone. Alternatively, the controller 104 can perform facial recognition. As other examples of techniques for obtaining the person's identification, the person can deliberately identify him or herself, such as offering a name tag for scanning, entering a PIN code or password, submitting biometric information (e.g., a fingerprint or retinal scan), speaking to allow for voice recognition. In another embodiment, the object-identification system 100 identifies the individual using skeletal tracking (i.e., the skeletal structure of the individual) and registers the skeletal structure. In addition to registering the person, the object-identification system 100 can record the person's time of arrival at the holding area.

At step 806, the object-identification system 100 associates the person with one or more objects in the holding area. The association can occur directly, from user input, or indirectly, based on an activity performed by the user and observed by the system. As an example of direct association, the system can expressly request that the person provide information about the purpose of the visit, such as depositing or removing an object, and the identity of each object the purpose involves. The person can provide this information through any number of input techniques, for example, scanning the label on a package to be deposited. Alternatively, the person can identify what the object is by typing in the name of the object or by speaking to the system, which uses voice recognition and speech-to-text conversion techniques. After receiving the information about each affected object, the system associates that object with the identity of the registered person.

As an example of indirect association, the object-identification system 100 can detect the activity performed by the person in the holding area. For example, through image processing, the system can detect that an object has been placed on or removed from a shelf and then associate the newly placed object, or the removed object, with the identity of the registered person.

At step 808, the object-identification system 100 attempts to recognize what the object is. Recognition may result from information supplied directly to the system by the user, for example, when the user enters that the “item is a microscope”; from a previous determination, for example, the system detects the removal of an object with an already known identity; or from object recognition, for example, the system executes its object recognition algorithm upon an image of the newly detected object. In one embodiment, the system automatically requests human interaction, namely, to ask the human to identify an object being deposited, moved, or removed. Such request can occur before, during, or after the system attempts its own object recognition.

A decision to request human interaction may be based on a threshold value derived by the controller 104 in its attempt at object recognition from a captured image. For example, if, at step 810, the threshold value exceeds a first (e.g., upper) threshold, the system considers an object to have been recognized with a high degree of confidence and may dispense with human interaction; if the threshold value is less than the first threshold but greater than a second (e.g., lower) threshold, the system considers an object to have been recognized, but with a moderate degree of confidence; if the threshold value falls below the second threshold, the system concludes it has not recognized any object in the image. The system may request that the person confirm or correct (step 812) the system’s identification if the determined threshold value is below the upper threshold, but above the lower threshold, and request (step 814) that the person provide the identification if the determined threshold value is below the lower threshold. Fewer or more than two thresholds may be used without departing from the principles described herein. Further, the system may request confirmation even if the threshold value exceeds the upper threshold or request the object’s identity in the event of an imprecise, incorrect, or unsuccessful object recognition.

The activity of the person in the holding area may change the layout of objects on the support surface. A new object has been placed, an object has been removed, an object has been moved to another location, or any combination thereof. The new arrangement of the objects produces different perspectives and varied angular irregularities in relation to the image and depth sensors of one or more sensor modules. Machine learning not only learns what an object looks like through both color and depth, it can now learn various perspectives on each object as they are placed in different locations in the area. This machine learning compensates for the dynamic perspectives of objects, seen by the image sensor, and learns how an identified object can be the same object if placed in different areas within the viewing area and at different angles, depths in the shelving. Accordingly, images now captured by the image sensors provide an opportunity to improve object recognition, through machine learning techniques. The system retrains the neural network with those newly captured images for which the neural network was unable to identify the object (at least initially) and needed the labeling information about the object provided by the user or by the neural network during a multi-pass authentication. The system can also record the person's time of departure when the person leaves the holding area, and then associate the person's time of arrival and time of departure with the object.

Consider the following illustrations as examples of operation of one embodiment of the object-identification system 100. Alice enters a room having several shelves. She is carrying a microscope and a smartphone. The smartphone is running Bluetooth®. The controller 104 connects to and communicates with the smartphone to establish the identity of the person as Alice. In addition, the controller establishes the time of Alice's entry into the room, for example, as 1:42 p.m., Thursday, April 16, 2019. Alice places the microscope on one of the shelves. Through image processing of images captured by the image sensor, the controller detects the object and location of the microscope. In addition, the controller may employ machine learning to recognize the object as a microscope. The controller may ask Alice to confirm its determination, whether the controller has recognized the object correctly or not. If the controller was unable to recognize the placed object, the controller may ask Alice to identify the object, which she may input electronically or verbally, depending upon the configuration of the object-identification system. Alternatively, the system may be configured to ask Alice the identity of the object, irrespective of its own recognition of the object. The system can then, locally or remotely on the server, immediately, or later, train its neural network with the images captured of the microscope and with the information, if any, provided by Alice. Alice then departs the room, and the controller records the time of departure as 1:48 p.m., Thursday, April 16, 2019.

Bob enters the room and submits his identification to the controller using a PIN code. The controller registers Bob and his time of entry as, for example, as 2:54 p.m., Thursday, April 16, 2019. The controller identifies Bob and, from his pattern of past practices, recognizes his regular use of the microscope. The controller asks, audibly or by a message displayed on a display screen, if Bob is looking for the microscope. If Bob answers in the affirmative, the controller illuminates the light

source and directs a light beam at the location on the shelves where the microscope resides. Bob removes the microscope from the shelf and departs the room with it. The system records Bob's time of departure as 2:56 p.m., Thursday, April 16, 2019, and that Bob has taken the microscope. By linking the arrival of the microscope with Alice, the removal of the microscope with Bob, the times of such operations, the presence of the microscope in the interim, all confirmed by video recordings, the system has thus established a chain of custody of the microscope. This chain of custody principle can extend to other fields of endeavor, such as processes for handling evidence. In the present context, chain of custody means a chronological recording of the sequence of custody (possession) and locations of physical objects coming into, moving within, and going out of the holding area. The object-identification system knows who has brought certain pieces of evidence into the evidence room, taken evidence from the room, and the precise locations of the evidence within the room in the interim, even if moved to another section within sight of the image sensor.

FIG. 9 shows an embodiment of a process 900 for recognizing objects in images. At step 902, the image-acquisition module 304 receives a captured image. Optionally, the image preprocessing module 305 edits (step 904) the image, including lowering its resolution, before sending (step 906) the image to the DNN 112. The DNN attempts to detect and identify (step 908) an object in the received image.

If the QMM 312 determines (step 910) that the DNN successfully identified one or more objects in the image, the object-identification system 100 uses (step 912) the information about each identified object, for example, for object-tracking purposes. The specific use of the object information depends upon the application for which the object-identification system is being used.

If, instead, the QMM determines (step 910) that the DNN was unsuccessful in the attempt to identify an object in the image, the AI module asks (step 914) the human to identify the object. After the human supplies the requested information, the optionally preprocessed image (produced in step 904) is stored (step 916) in the image database 122 with the human-provided labeling information, for later use in retraining the DNN.

In one embodiment, shown in phantom in FIG. 9, the AI module performs (step 918) the multi-pass authentication process described in connection with FIG. 7. If, at step 920, the DNN is unable to identify an object in the second pass, the AI module asks (step 914) the human to identify the object, and the optionally preprocessed image is stored (step 916) in the image database 122 with the human-provided labeling information. Alternatively, if, at step 920, the DNN successfully identified an object in the image on the second attempt, the optionally preprocessed image is stored (step 916) with DNN-provided labeling information for later use in retraining the DNN. Accordingly, when the DNN can identify an object on the second pass, the image is saved with the identifying label produced by the DNN; when the DNN is unable to identify an object on the second attempt, the image is stored with the human-supplied label information.

FIG. 10 shows an embodiment of a process 1000 for retraining the DNN with images in which the DNN does not initially identify an object. Periodically, the object-identification system 100 determines that the DNN is to be retrained with images stored in the image database. Such retraining can be based on a schedule (e.g., every evening, every week, etc.).

5 For purposes of retraining the DNN 112 (FIG. 1), the DNN trainer 314 (e.g., FIG. 3) retrieves (step 1002) the images in the image database 122. Such images are stored in local storage 124 or remote storage 128. The image database 122 contains the initial training set of images and each image in which the DNN 112 was initially unable to identify an object.

10 Based on the images in the image database, the DNN trainer 314 runs program code for neural network training 222 (FIG. 2) that produces (step 1004) a new neural network weight file. The new weight file contains a weighting (i.e., parameter) value for each of the “neurons” of the DNN 112. To update the DNN, that is, to retrain the DNN, the DNN trainer sends this new weight file to the DNN, to be applied by the DNN to its neuron structure for object detection and recognition in subsequently received images.

15 In general, the DNN trainer maintains a copy of the current weight file for the DNN. The retraining of the DNN can occur in whole or in part. When retraining in whole, the entire DNN is trained from scratch, that is, the current weight file is erased and replaced with a newly generated weight file. It’s as though the DNN was again a blank slate and was being initially trained. This retraining uses the initial training set of images and each additional image added to the image
20 database for not being initially identified.

When retraining in part, the retraining can focus on certain layers of the DNN. For example, consider a DNN with ten hidden layers; retraining can be performed on the seventh, eighth, and ninth hidden layers only, the operative principle being to avoid performing a full DNN training, which can be time consuming, when a focused retraining can suffice. In this example, only those parameter
25 values in the current weight file that are associated with the neurons of the seventh, eighth, and ninth hidden layers are changed. The new weight file, produced by the DNN trainer and sent to the DNN, is a mix of the new parameter values for the neurons of the seventh, eighth, and ninth hidden layers and old parameter values for the remainder of the DNN.

30 Consider that, as an example of the machine learning in operation, Alice places a microscope on a shelf. At the time of registering Alice, if the object-identification system 100 does not recognize the object, the system asks Alice to identify the placed object; she may respond that the object is a microscope. The object-identification system 100 further captures one or more images of the object on the shelf and associates each captured image with the information provided by Alice (i.e., the object is a microscope). The DNN trainer uses each captured image and the information provided by
35 Alice to train the neural network 112. This training may be the system’s initial training for identifying microscopes, or cumulative to the system’s present capability. In either case, after the training, the object-identification system is better suited for identifying microscopes.

Although described with respect to detecting, tracking, and recognizing objects, the machine-learning techniques described herein extend to the detecting, tracking, and recognizing faces, skeletal structure, body position, and movement of people in the captured images. In similar fashion as images of objects are used to train the deep neural networks to improve object recognition, images of faces can be used to train such networks to improve facial recognition for purposes of user registration, and images of skeletal features, such as hands, arms, and legs can be used to train such networks to improve for purposes of identifying and tracking individual persons and objects.

As will be appreciated by one skilled in the art, aspects of the systems described herein may be embodied as a system, method, and computer program product. Thus, aspects of the systems described herein may be embodied in entirely hardware, in entirely software (including, but not limited to, firmware, program code, resident software, microcode), or in a combination of hardware and software. All such embodiments may generally be referred to herein as a circuit, a module, or a system. In addition, aspects of the systems described herein may be in the form of a computer program product embodied in one or more computer readable media having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. The computer readable medium may be a non-transitory computer readable storage medium, examples of which include, but are not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination thereof.

As used herein, a computer readable storage medium may be any tangible medium that can contain or store a program for use by or in connection with an instruction execution system, apparatus, device, computer, computing system, computer system, or any programmable machine or device that inputs, processes, and outputs instructions, commands, or data. A non-exhaustive list of specific examples of a computer readable storage medium include an electrical connection having one or more wires, a portable computer diskette, a floppy disk, a hard disk, a random access memory (RAM), a read-only memory (ROM), a USB flash drive, an non-volatile RAM (NVRAM or NOVRAM), an erasable programmable read-only memory (EPROM or Flash memory), a flash memory card, an electrically erasable programmable read-only memory (EEPROM), an optical fiber, a portable compact disc read-only memory (CD-ROM), a DVD-ROM, an optical storage device, a magnetic storage device, or any suitable combination thereof.

A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction

execution system, apparatus, or device. As used herein, a computer readable storage medium is not a computer readable propagating signal medium or a propagated signal.

Program code may be embodied as computer-readable instructions stored on or in a computer readable storage medium as, for example, source code, object code, interpretive code, executable
5 code, or combinations thereof. Any standard or proprietary, programming or interpretive language can be used to produce the computer-executable instructions. Examples of such languages include Python, C, C++, Pascal, JAVA, BASIC, Smalltalk, Visual Basic, and Visual C++.

Transmission of program code embodied on a computer readable medium can occur using any appropriate medium including, but not limited to, wireless, wired, optical fiber cable, radio frequency
10 (RF), or any suitable combination thereof.

The program code may execute entirely on a user's device, partly on the user's device, as a stand-alone software package, partly on the user's device and partly on a remote computer or entirely on a remote computer or server. Any such remote computer may be connected to the user's device through any type of network, including a local area network (LAN) or a wide area network (WAN), or the
15 connection may be made to an external computer (for example, through the Internet using an ISP (Internet Service Provider)).

Additionally, the methods described herein can be implemented on a special purpose computer, a programmed microprocessor or microcontroller and peripheral integrated circuit element(s), an ASIC or other integrated circuit, a digital signal processor, a hard-wired electronic or logic circuit such as
20 discrete element circuit, a programmable logic device such as PLD, PLA, FPGA, PAL, or the like. In general, any device capable of implementing a state machine that is in turn capable of implementing the proposed methods herein can be used to implement the principles described herein.

Furthermore, the disclosed methods may be readily implemented in software using object or object-oriented software development environments that provide portable source code that can be
25 used on a variety of computer or workstation platforms. Alternatively, the disclosed system may be implemented partially or fully in hardware using standard logic circuits or a VLSI design. Whether software or hardware is used to implement the systems in accordance with the principles described herein is dependent on the speed and/or efficiency requirements of the system, the particular function, and the particular software or hardware systems or microprocessor or microcomputer systems being
30 utilized. The methods illustrated herein however can be readily implemented in hardware and/or software using any known or later developed systems or structures, devices and/or software by those of ordinary skill in the applicable art from the functional description provided herein and with a general basic knowledge of the computer and image processing arts.

Moreover, the disclosed methods may be readily implemented in software executed on
35 programmed general-purpose computer, a special purpose computer, a microprocessor, or the like. In these instances, the systems and methods of the principles described herein may be implemented as program embedded on personal computer such as JAVA® or CGI script, as a resource residing on a

server or graphics workstation, as a plug-in, or the like. The system may also be implemented by physically incorporating the system and method into a software and/or hardware system.

While the aforementioned principles have been described in conjunction with a number of embodiments, it is evident that many alternatives, modifications, and variations would be or are
5 apparent to those of ordinary skill in the applicable arts. References to "one embodiment" or "an embodiment" or "another embodiment" means that a particular, feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment described herein. A reference to a particular embodiment within the specification do not necessarily all refer to the
10 same embodiment. The features illustrated or described in connection with one exemplary embodiment may be combined with the features of other embodiments. Accordingly, it is intended to embrace all such alternatives, modifications, equivalents, and variations that are within the spirit and scope of the principles described herein.

What is claimed is:

CLAIMS:

1. An object-identification system comprising:
 - an image sensor configured to capture images of objects disposed in an area
5 designated for holding objects;
 - a deep neural network trained to detect and recognize objects in images like those objects held in the designated area; and
 - a controller in communication with the image sensor to receive images captured by the image sensor and with the deep neural network, the controller including one or more
10 processors configured to register an identity of a person who visits the area designated for holding objects, to submit an image to the deep neural network, to associate the registered identity of the person with an object detected in the image submitted to the deep neural network, to retrain the deep neural network using the submitted image if the deep neural network is unable to recognize the object detected in the submitted image, and to track a
15 location of the detected object while the detected object is in the area designated for holding objects.
2. The object-identification system of claim 1, wherein the controller is further configured to acquire labeling information for the detected object in response to the deep neural network being unable to recognize the detected object in the submitted image, to associate the labeling
20 information with the version of the image submitted to the deep neural network, and to store the version of the image and associated labeling information in an image database used to retrain the deep neural network.
3. The object-identification system of claim 2, further comprising a human-input acquisition module configured to acquire the labeling information from a user in response to a request
25 from the controller when the deep neural network is unable to recognize the detected object in the image in the submitted image.
4. The object-identification system of claim 1, wherein the controller is further configured to find an area within the image in which a change appears if the deep neural network is unable to recognize the detected object, to produce a version of the image that focuses upon the area
30 of change, and to submit the version of the image to the deep neural network to determine whether the deep neural network is able to recognize the detected object in the second version of the image.
5. The object-identification system of claim 4, wherein the controller is further configured to acquire labeling information for the detected object irrespective of whether the deep neural
35 network recognizes the detected object in the submitted version of the image, to associate the acquired labeling information with the version of the image submitted to the deep neural

network, and to store the version of the image and associated labeling information in an image database used to retrain the deep neural network.

6. The object-identification system of claim 5, wherein the controller is further configured to acquire the labeling information from the deep neural network when the deep neural network recognizes the detected object in the submitted version of the image.
7. The object-identification system of claim 1, wherein the deep neural network is a first deep neural network, and further comprising a second deep neural network configured to operate in parallel to the first deep neural network, each of the first and second deep neural networks producing an output based on image data obtained from the image, wherein the image data obtained by the first deep neural network are different from the image data obtained by the second deep neural network.
8. The object-identification system of claim 1, further comprising a depth sensor with a field of view that substantially matches a field of view of the image sensor, the depth sensor acquiring depth pixels value of images within its field of view, wherein a depth pixel value and less than three pixel values taken from the group consisting of R (red), G (green), and B (blue) are submitted as image data to the deep neural network when the image is submitted to the deep neural network during training or object recognition.
9. The object-identification system of claim 1, wherein the deep neural network resides on a remote server system, and the controller further comprises a network interface to communicate with the deep neural network on the server system.
10. A method of identifying and tracking objects, the method comprising the steps of:
 - registering an identity of a person who visits an area designated for holding objects;
 - capturing an image of the area designated for holding objects;
 - submitting a version of the image to a deep neural network trained to detect and recognize objects in images like those objects held in the designated area;
 - detecting an object in the version of the image;
 - associating the registered identity of the person with the detected object;
 - retraining the deep neural network using the version of the image if the deep neural network is unable to recognize the detected object; and
 - tracking a location of the detected object while the detected object is in the area designated for holding objects.
11. The method of claim 10, further comprising:
 - acquiring labeling information for the object detected in the version of the image in response to the deep neural network being unable to recognize the detected object in the version of the image;
 - associating the labeling information with the version of the image; and

storing the version of the captured image and associated labeling information in an image database used to retrain the deep neural network.

12. The method of claim 11, wherein the step of acquiring labeling information for the object detected in the version of the image in response to the deep neural network being unable to recognize the detected object in the version of the image comprises acquiring the labeling information from user-supplied input.

13. The method of claim 10, wherein the version of the image is a first version of the image, and further comprising:

finding an area within the first version of the image in which a change appears when the deep neural network is unable to recognize the object detected in the first version of the image;

producing a second version of the image that focuses upon the found area of change; and

submitting the second version of the image to the deep neural network to determine whether the deep neural network can recognize the detected object in the second version of the image.

14. The method of claim 13, further comprising:

acquiring labeling information for the object detected in the first version of the image irrespective of whether the deep neural network recognizes the detected object in the second version of the image;

associating the labeling information with the first version of the image; and

storing the first version of the captured image and associated labeling information in an image database used to retrain the deep neural network.

15. The method of claim 14, wherein the step of acquiring labeling information for the object detected in the version of the image comprises acquiring the labeling information from the deep neural network when the deep neural network recognizes the detected object in the version of the image.

16. The method of claim 10, wherein the deep neural network is a first deep neural network, and further comprising the step of submitting image data, acquired from the version of the image, to the first deep neural network and a second deep neural network in parallel, wherein the image data submitted to the first deep neural network are different from the image data submitted to the second deep neural network.

17. The method of claim 10, wherein the step of submitting a version of the image to the deep neural network comprises submitting a depth pixel value and less than three pixel values taken from the group consisting of R (red), G (green), and B (blue) as image data to the deep neural network.

18. A sensor module comprising:

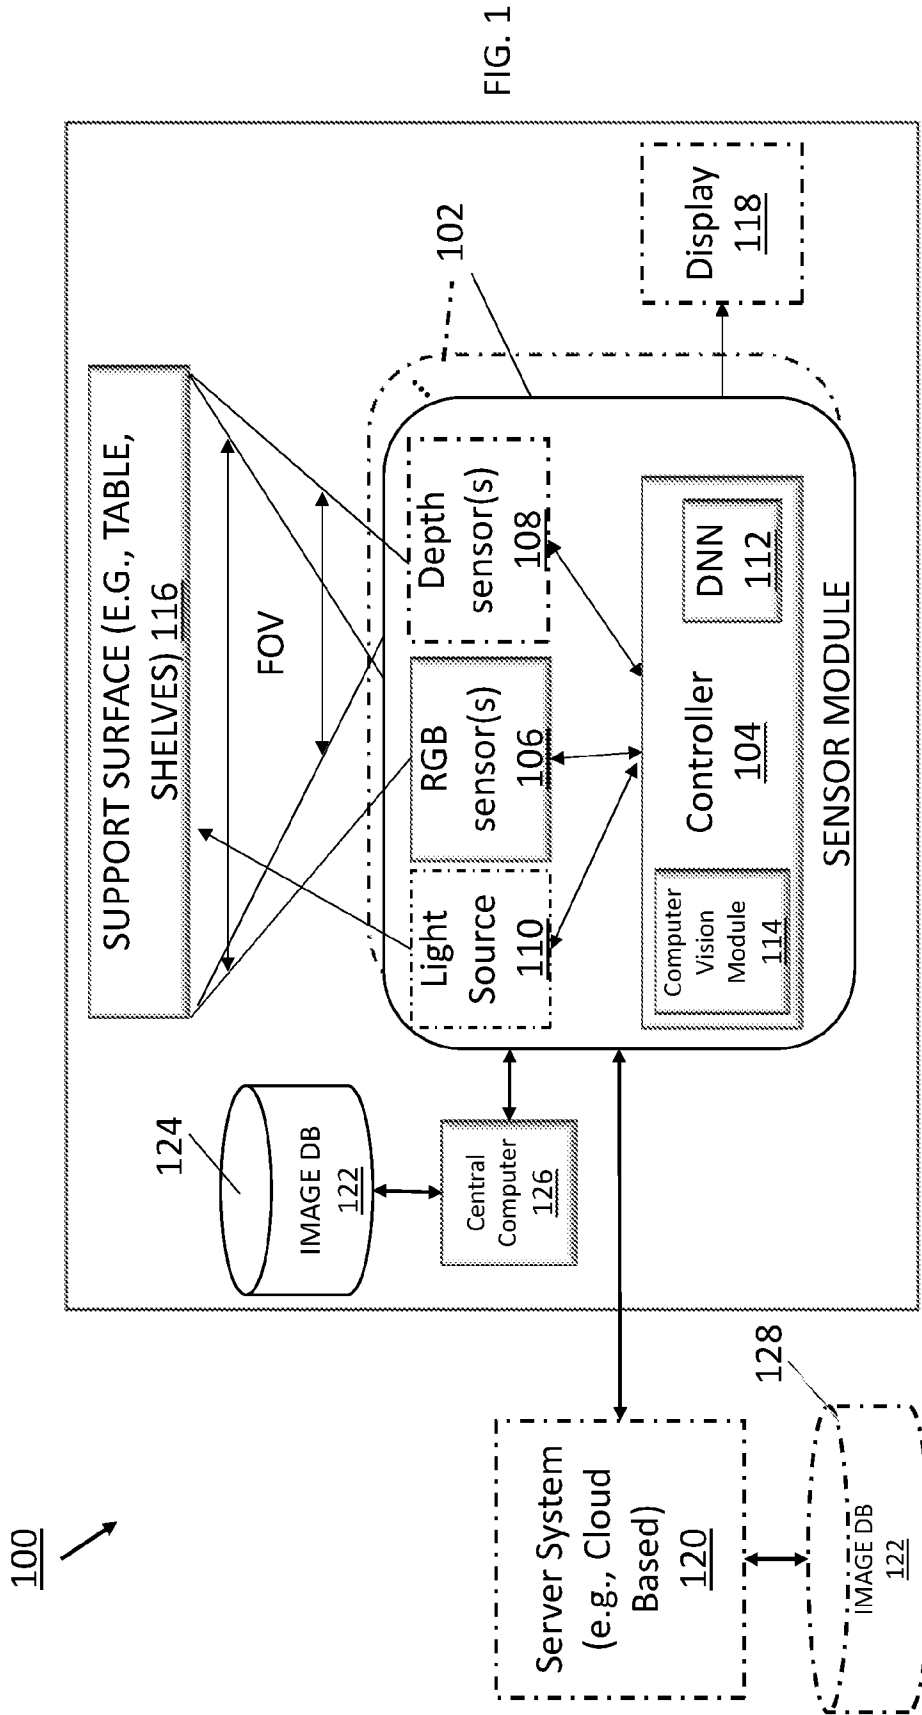
an image sensor configured to capture an image within its field of view;

a depth sensor having a field of view that substantially matches the field of view of the image sensor, the depth sensor being configured to acquire estimated depth values for an image captured by the depth sensor; and

a controller in communication with the image sensor and depth sensor to receive image data associated with the image captured by the image sensor and estimated depth values associated with the image captured by the depth sensor, the controller including one or more processors configured to register an identity of a person who visits an area designated for holding objects, to submit the image data associated with the image captured by the image sensor and the estimated depth values associated with the image captured by the depth sensor to a deep neural network trained to detect and recognize objects in images like those objects held in the designated area, to associate the registered identity of the person with an object detected in the image data and estimated depth values submitted to the deep neural network, and to save a version of the images captured by the image sensor and the depth sensor for use in subsequent retraining of the deep neural network if the deep neural network is unable to recognize the detected object.

19. The sensor module of claim 18, wherein the controller further comprises a cloud interface to communicate with the deep neural network over a network.

20. The sensor module of claim 18, further comprising a human-input acquisition module configured to acquire the labeling information from a user in response to a request from the controller when the deep neural network is unable to recognize the detected object based on the submitted image data and estimated depth values.



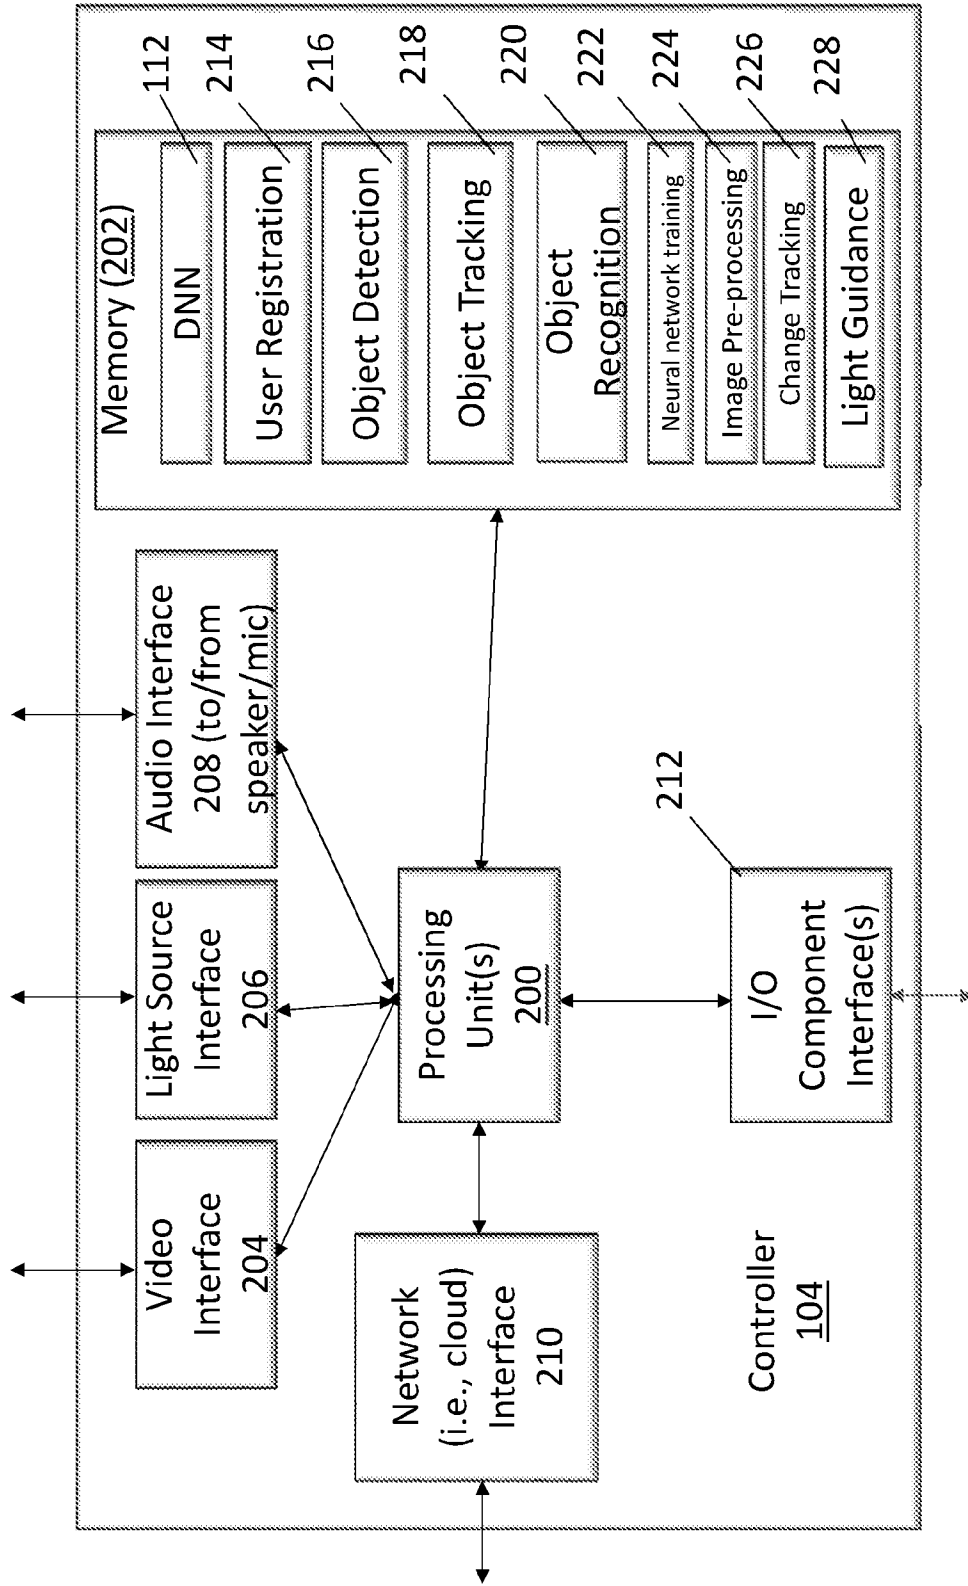


FIG. 2




FIG. 3

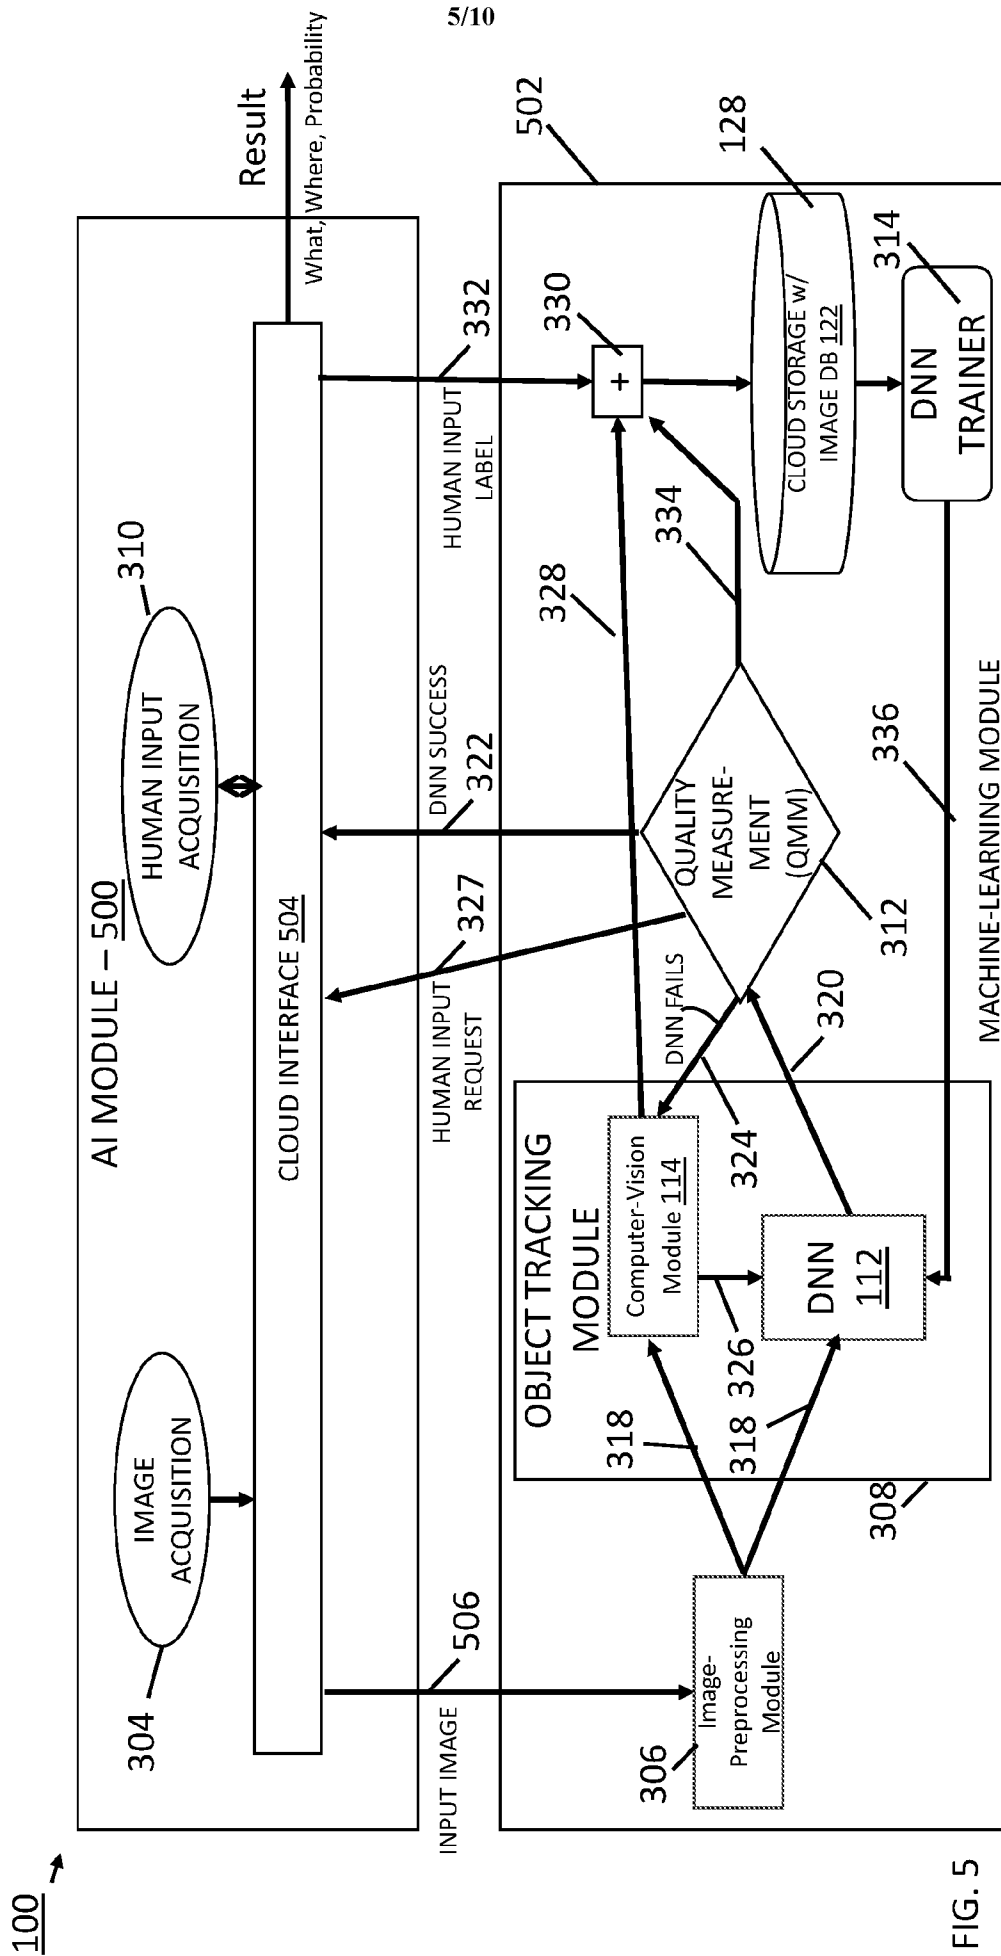


FIG. 5

100 →




FIG. 6




FIG. 7

800

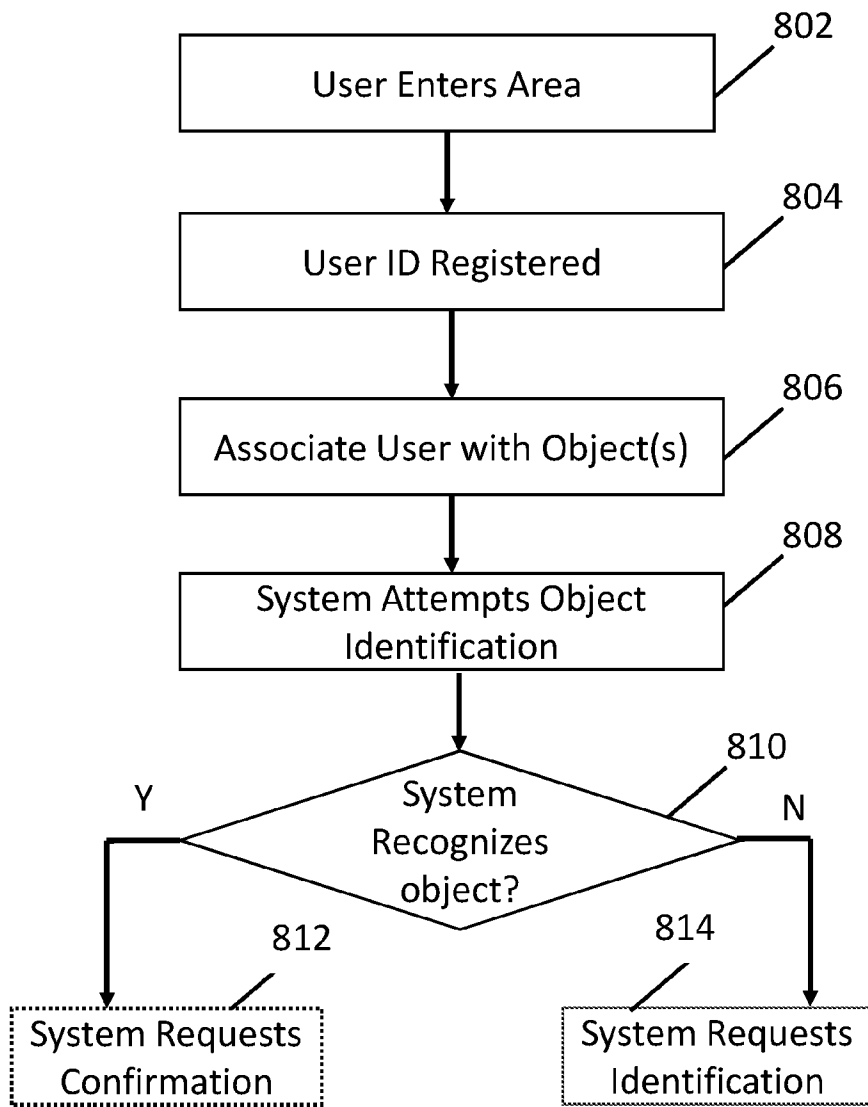


FIG. 8

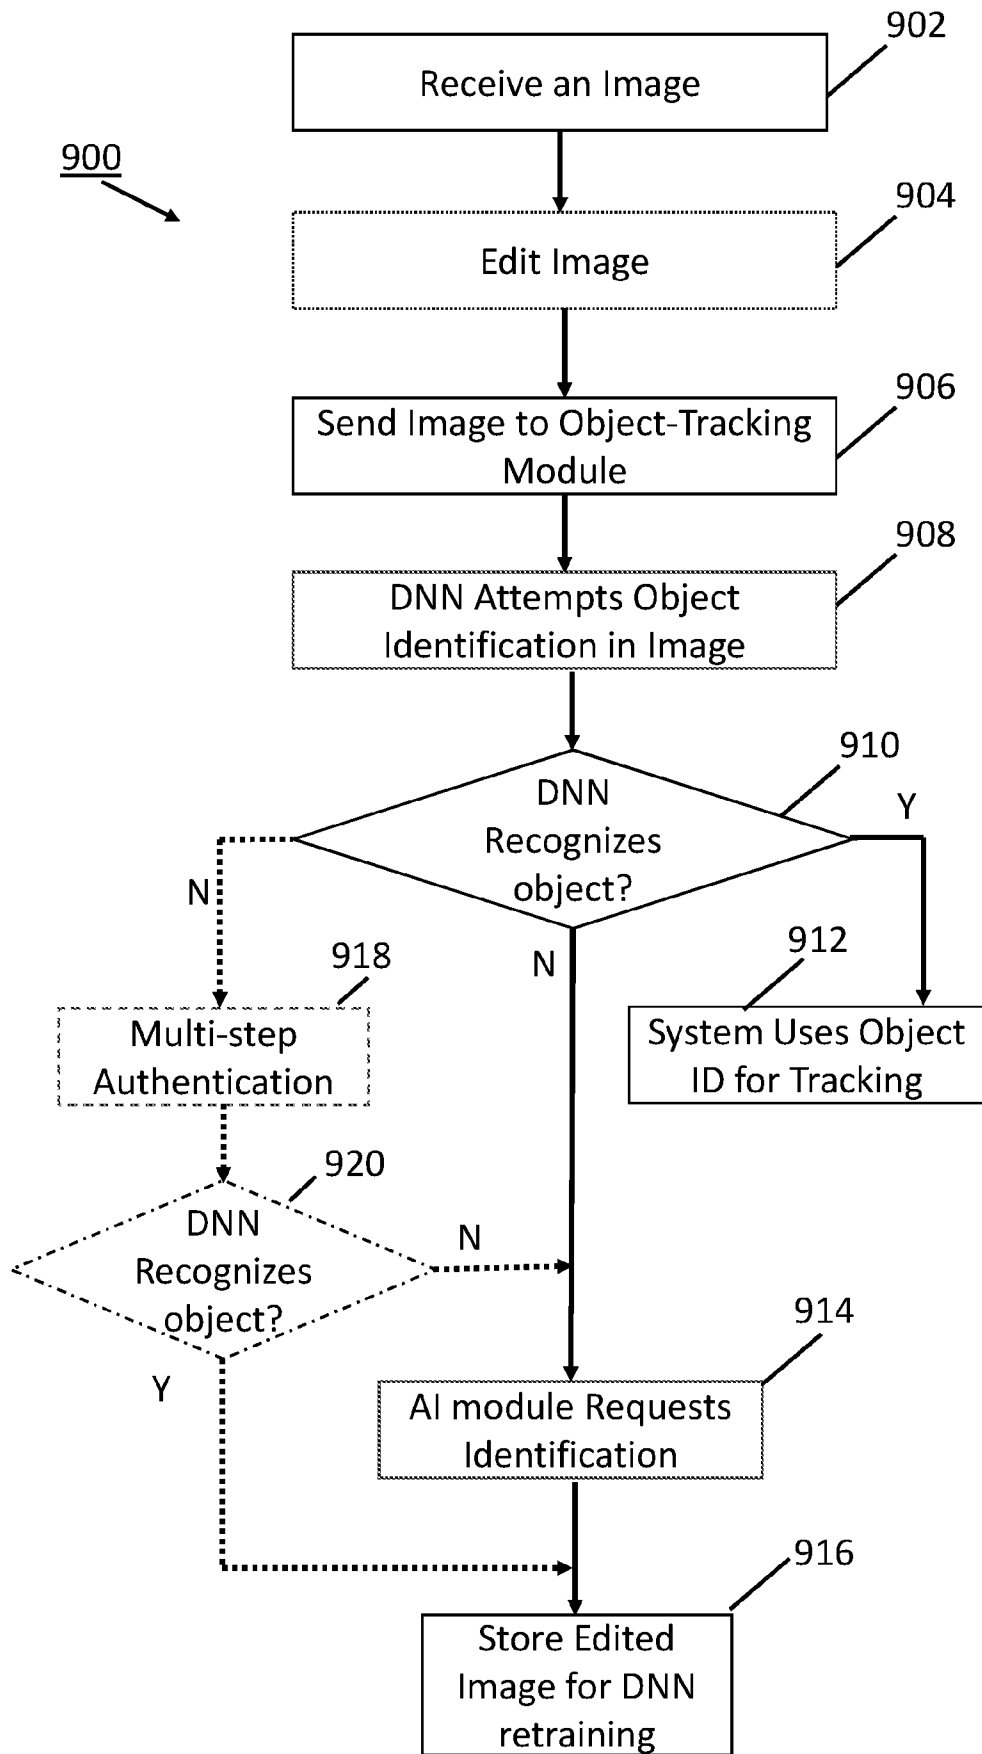


FIG. 9

1000
↓




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 19/51874

A. CLASSIFICATION OF SUBJECT MATTER

IPC - G06K 9/66 (2019.01)

CPC - G06K 9/00288; G06K 9/00624; G06K9/66; G06N 3/0454

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 6,492,905 B2 (MATHIAS et al.) 10 December 2002 (10.12.2002), Fig 1-3, abstract, col 3, ln 3-col 4, ln 59, col 6, ln 59-col 8, ln 17, col 10, ln 30-61	1-20
A	US 9,594,983 B2 (DIGIMARC CORPORATION) 14 March 2017 (14.03.2017), abstract, col 6, ln 39-54, col 7, ln 6-27, col 24, ln 1-6	1-20
A	US 9,740,937 B2 (AVIGILON FORTRESS CORPORATION) 22 August 2017 (22.08.2017), Fig 1, 2, abstract, col 6, ln 43-61	1-20
A	US 2018/0197218 A1 (VERIZON PATENT AND LICENSING INC.) 12 July 2018 (12.07.2018), Fig 1, 6B, abstract, para [0014], [0063], [0068]	1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"D" document cited by the applicant in the international application

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

21 November 2019

Date of mailing of the international search report

13 DEC 2019

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Lee Young

Telephone No. PCT Helpdesk: 571-272-4300