(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
02 July 2020 (02.07.2020)

WIPO | PCT

(10) International Publication Number
WO 2020/139478 A1

(51) International Patent Classification:
G06T 7/73 (2017.01)

(21) International Application Number:
PCT/US2019/061955

(22) International Filing Date:
18 November 2019 (18.11.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/232,147 26 December 2018 (26.12.2018) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Blvd., Santa Clara, CA 95054 (US).

(72) Inventors: **KAVULYA, Soila**; 5282 SE Thornapple Street, Hillsboro, OR 97123 (US). **CHATTOPADHYAY, Rita**; 4700 W. Geronimo Street, Chandler, AZ 85226 (US). **MARTINEZ-CANALES, Monica, Lucia**; 1237 Woodview Terrace, Los Altos, CA 94024 (US).

(74) Agent: **KELLETT, Glen, M.**; Barnes & Thornburg LLP, 11 South Meridian Street, Indianapolis, IN 46204 (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,

(54) **Title:** TECHNOLOGIES FOR FUSING DATA FROM MULTIPLE SENSORS TO IMPROVE OBJECT DETECTION, IDENTIFICATION, AND LOCALIZATION

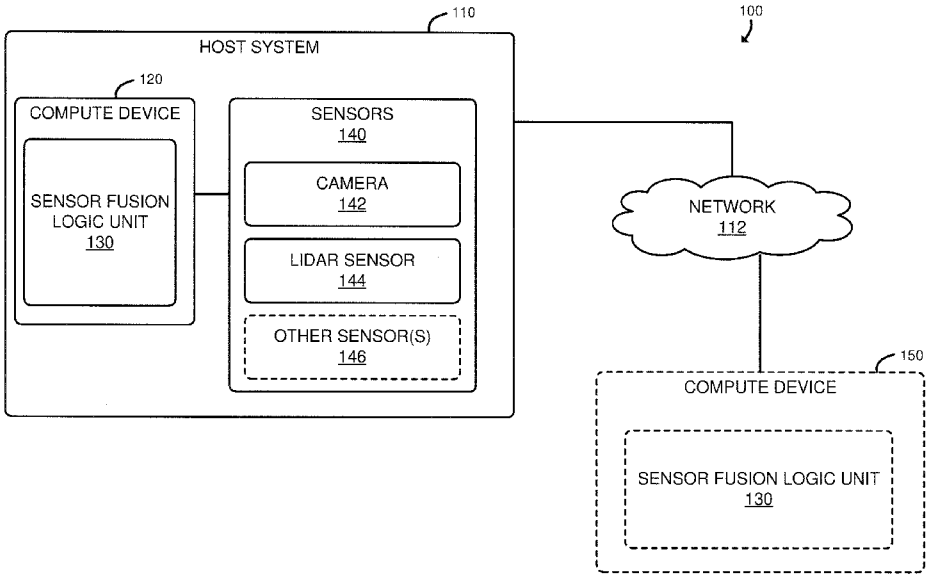

FIG. 1

(57) **Abstract:** Technologies for performing sensor fusion include a compute device. The compute device includes circuitry configured to obtain detection data indicative of objects detected by each of multiple sensors of a host system. The detection data includes camera detection data indicative of a two or three dimensional image of detected objects and lidar detection data indicative of depths of detected objects. The circuitry is also configured to merge the detection data from the multiple sensors to define final bounding shapes for the objects.

WO 2020/139478 A1

[Continued on next page]

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

TECHNOLOGIES FOR FUSING DATA FROM MULTIPLE SENSORS TO IMPROVE OBJECT DETECTION, IDENTIFICATION, AND LOCALIZATION

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application claims the benefit of U.S. Patent Application Serial No. 16/232147, filed December 26, 2018.

BACKGROUND

[0002] Efficient and robust sensing of objects in an environment is of high importance to automated systems such as robots, drones, and autonomous vehicles. State of the art systems typically include three functional processes: detection of an object (e.g., determining that an object exists), recognition (e.g., determining what the object is), and localization (e.g., determining where the object is). In performing these operations, state of the art systems typically exhibit several shortcomings, including incorrect reconciliation of detections based on rejection of bounding box proposals from multiple sensors, detection but not recognition of objects within cells, and incorrect localization based on conducting the merging of detection data in two dimensional space.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The concepts described herein are illustrated by way of example and not by way of limitation in the accompanying figures. For simplicity and clarity of illustration, elements illustrated in the figures are not necessarily drawn to scale. Where considered appropriate, reference labels have been repeated among the figures to indicate corresponding or analogous elements.

[0004] FIG. 1 is a simplified diagram of at least one embodiment of a system for fusing data from multiple sensors for object detection, identification, and localization;

[0005] FIG. 2 is a simplified block diagram of at least one embodiment of a compute device included in the system of FIG. 1;

[0006] FIGS. 3-5 are simplified block diagrams of at least one embodiment of a method for fusing data from multiple sensors that may be performed by a compute device of the system of FIG. 1;

[0007] FIG. 6 is a simplified diagram of at least one embodiment of a pipeline for performing sensor fusion with the system of FIG. 1; and

[0008] FIG. 7 is a chart of at least one embodiment of a set of relative weights that may be applied by the system of FIG. 1 to sensor data when performing sensor fusion.

DETAILED DESCRIPTION

[0009] While the concepts of the present disclosure are susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and will be described herein in detail. It should be understood, however, that there is no intent to limit the concepts of the present disclosure to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives consistent with the present disclosure and the appended claims.

[0010] References in the specification to "one embodiment," "an embodiment," "an illustrative embodiment," etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may or may not necessarily include that particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described. Additionally, it should be appreciated that items included in a list in the form of "at least one A, B, and C" can mean (A); (B); (C); (A and B); (A and C); (B and C); or (A, B, and C). Similarly, items listed in the form of "at least one of A, B, or C" can mean (A); (B); (C); (A and B); (A and C); (B and C); or (A, B, and C).

[0011] The disclosed embodiments may be implemented, in some cases, in hardware, firmware, software, or any combination thereof. The disclosed embodiments may also be implemented as instructions carried by or stored on a transitory or non-transitory machine-readable (e.g., computer-readable) storage medium, which may be read and executed by one or more processors. A machine-readable storage medium may be embodied as any storage device, mechanism, or other physical structure for storing or transmitting information in a form readable by a machine (e.g., a volatile or non-volatile memory, a media disc, or other media device).

[0012] In the drawings, some structural or method features may be shown in specific arrangements and/or orderings. However, it should be appreciated that such specific arrangements and/or orderings may not be required. Rather, in some embodiments, such features may be arranged in a different manner and/or order than shown in the illustrative figures. Additionally, the inclusion of a structural or method feature in a particular figure is not meant to imply that such feature is required in all embodiments and, in some embodiments, may not be included or may be combined with other features.

[0013] Referring now to FIG. 1, a system 100 for performing sensor fusion includes a host system 110, which may be embodied as any device or assembly of devices (e.g., an

autonomous vehicle, a drone, a robot, etc.) configured to identify objects in an environment around the host system 110 and potentially respond to the presence of the objects (e.g., avoid a collision with a pedestrian crossing a street, stop at a stop sign, etc.). In the illustrative embodiment, the host system 110 includes a compute device 120 that includes a sensor fusion logic unit 130 which may be embodied as any device or circuitry (e.g., a processor, a field programmable gate array (FPGA), an application specific integrated circuit (ASIC), etc.) configured to receive detection data from a set of sensors 140 of the host system 110 and merge the detection data to define bounding shapes (e.g., shapes having boundaries that enclose an object) detected by the different sensors 140. In the illustrative embodiment, the sensors 140 include a camera 142, which may be embodied as any device or circuitry (e.g., one or more lenses adapted to direct light into an electronic sensor, such as a charge-coupled device) configured to convert light (e.g., light in a visible spectrum) to pixel information (e.g., sets of color values for each of multiple x and y coordinates in a two dimensional frame, or x, y and z coordinates in a three dimensional frame using 3D cameras such as Intel RealSense™ Depth cameras). Additionally, in the illustrative embodiment, the sensors 140 include a lidar sensor 144 which may be embodied as any device or circuitry configured to illuminate a target objects with pulsed laser light and measure reflected pulses from the target objects to determine the positions of the target objects relative to the lidar sensor 144 in a three dimensional coordinate system (e.g., x, y, and z coordinates). As such, the lidar sensor 144 provides depth information, which may not be provided by the camera 142.

[0014] The sensors 140 may also include one or more other sensors 146 (e.g., a radar sensor to determine the range of objects using radio waves, an acoustic sensor to detect sounds in the environment, a multispectral or hyperspectral image sensor to differentiate between types of objects using their electromagnetic spectrum, etc.). In performing the sensor fusion, the sensor fusion logic unit 130, in the illustrative embodiment, selectively applies different weights to the data received from the sensors 140 (e.g., the camera 142 and lidar sensor 144) as a function of a context in which the data was obtained (e.g., weather, amount of available light in the environment, etc.) and known attributes of the corresponding sensors 140 (e.g., the camera 142 is less reliable in relatively dark environments compared to the lidar sensor 144). Furthermore, the sensor fusion logic unit 130 applies one or more rules, which may be user-defined or otherwise configurable (e.g., in updates from a manufacturer), that produce degrees of belief that object detections represented in the data from the various sensors 140 represent the same object, different objects, or unoccupied space ("free space"). Further, the sensor fusion logic unit 130 determines whether to merge (e.g., fuse) bounding boxes representing objects in the data from different sensors 140 to define a resulting bounding shape or to leave certain bounding boxes

separate. In some embodiments, a compute device 150 located remotely from the host system 110 may receive the sensor data (e.g., through a network 112), perform the sensor fusion operations, and provide data indicative of the bounding shapes to the host system 110 (e.g., to the compute device 120) for analysis. As compared to state of the art systems, which may (i) perform non-maximum suppression (e.g., keep one overlapping bounding box and discard the rest of the overlapping bounding boxes), (ii) execute occupancy-grid object detection (e.g., determine whether each cell in a grid is occupied by an object, but potentially fail to identify distinct objects in an environment), and/or (iii) attempt to merge detected objects based only on two dimensional information (e.g., without depth information), the system 100 is more efficient and reliable at performing object detection, identification, and localization.

[0015] Referring now to FIG. 2, the illustrative compute device 120 includes a compute engine (also referred to herein as "compute engine circuitry") 210, an input/output (I/O) subsystem 216, communication circuitry 218, and one or more data storage devices 222. Of course, in other embodiments, the compute device 120 may include other or additional components, such as those commonly found in a computer (e.g., a display, peripheral devices, etc.). Additionally, in some embodiments, one or more of the illustrative components may be incorporated in, or otherwise form a portion of, another component. The compute engine 210 may be embodied as any type of device or collection of devices capable of performing various compute functions described below. In some embodiments, the compute engine 210 may be embodied as a single device such as an integrated circuit, an embedded system, a field-programmable gate array (FPGA), a system-on-a-chip (SOC), or other integrated system or device. In the illustrative embodiment, the compute engine 210 includes or is embodied as a processor 212 and a memory 214. The processor 212 may be embodied as any type of processor capable of performing the functions described herein. For example, the processor 212 may be embodied as a multi-core processor(s), a microcontroller, or other processor or processing/controlling circuit. In some embodiments, the processor 212 may be embodied as, include, or be coupled to an FPGA, an application specific integrated circuit (ASIC), reconfigurable hardware or hardware circuitry, or other specialized hardware to facilitate performance of the functions described herein. In the illustrative embodiment, the processor 212 includes the sensor fusion logic unit 130, described with reference to FIG. 1. However, in other embodiments, the sensor fusion logic unit 130 may be located elsewhere in the compute device 120 (e.g., not integrated into the processor 212).

[0016] The main memory 214 may be embodied as any type of volatile (e.g., dynamic random access memory (DRAM), etc.) or non-volatile memory or data storage capable of performing the functions described herein. Volatile memory may be a storage medium that

requires power to maintain the state of data stored by the medium. Non-limiting examples of volatile memory may include various types of random access memory (RAM), such as dynamic random access memory (DRAM) or static random access memory (SRAM). One particular type of DRAM that may be used in a memory module is synchronous dynamic random access memory (SDRAM). In particular embodiments, DRAM of a memory component may comply with a standard promulgated by JEDEC, such as JESD79F for DDR SDRAM, JESD79-2F for DDR2 SDRAM, JESD79-3F for DDR3 SDRAM, JESD79-4A for DDR4 SDRAM, JESD209 for Low Power DDR (LPDDR), JESD209-2 for LPDDR2, JESD209-3 for LPDDR3, and JESD209-4 for LPDDR4. Such standards (and similar standards) may be referred to as DDR-based standards and communication interfaces of the storage devices that implement such standards may be referred to as DDR-based interfaces.

[0017] In one embodiment, the memory device is a block addressable memory device, such as those based on NAND or NOR technologies. A memory device may also include a three dimensional crosspoint memory device (e.g., Intel 3D XPoint™ memory), or other byte addressable write-in-place nonvolatile memory devices. In one embodiment, the memory device may be or may include memory devices that use chalcogenide glass, multi-threshold level NAND flash memory, NOR flash memory, single or multi-level Phase Change Memory (PCM), a resistive memory, nanowire memory, ferroelectric transistor random access memory (FeTRAM), anti-ferroelectric memory, magnetoresistive random access memory (MRAM) memory that incorporates memristor technology, resistive memory including the metal oxide base, the oxygen vacancy base and the conductive bridge Random Access Memory (CB-RAM), or spin transfer torque (STT)-MRAM, a spintronic magnetic junction memory based device, a magnetic tunneling junction (MTJ) based device, a DW (Domain Wall) and SOT (Spin Orbit Transfer) based device, a thyristor based memory device, or a combination of any of the above, or other memory. The memory device may refer to the die itself and/or to a packaged memory product.

[0018] In some embodiments, 3D crosspoint memory (e.g., Intel 3D XPoint™ memory) may comprise a transistor-less stackable cross point architecture in which memory cells sit at the intersection of word lines and bit lines and are individually addressable and in which bit storage is based on a change in bulk resistance. In some embodiments, all or a portion of the main memory 214 may be integrated into the processor 212. In operation, the main memory 214 may store various software and data used during operation such as applications, libraries, and drivers.

[0019] The compute engine 210 is communicatively coupled to other components of the compute device 120 via the I/O subsystem 216, which may be embodied as circuitry and/or components to facilitate input/output operations with the compute engine 210 (e.g., with the processor 212 and/or the main memory 214) and other components of the compute device 120

and other devices (e.g., the sensors 140). For example, the I/O subsystem 216 may be embodied as, or otherwise include, memory controller hubs, input/output control hubs, integrated sensor hubs, firmware devices, communication links (e.g., point-to-point links, bus links, wires, cables, light guides, printed circuit board traces, etc.), and/or other components and subsystems to facilitate the input/output operations. In some embodiments, the I/O subsystem 216 may form a portion of a system-on-a-chip (SoC) and be incorporated, along with one or more of the processor 212, the main memory 214, and other components of the compute device 120, into the compute engine 210.

[0020] The communication circuitry 218 may be embodied as any communication circuit, device, or collection thereof, capable of enabling communications over the network 112 between the compute device 120 and another device (e.g., the compute device 150, etc.). The communication circuitry 218 may be configured to use any one or more communication technology (e.g., wired or wireless communications) and associated protocols (e.g., Ethernet, Bluetooth®, Wi-Fi®, WiMAX, etc.) to effect such communication.

[0021] The illustrative communication circuitry 218 includes a network interface controller (NIC) 220, which may also be referred to as a host fabric interface (HFI). The NIC 220 may be embodied as one or more add-in-boards, daughter cards, network interface cards, controller chips, chipsets, or other devices that may be used by the compute device 120 to connect with another compute device (e.g., the compute device 150, etc.). In some embodiments, the NIC 220 may be embodied as part of a system-on-a-chip (SoC) that includes one or more processors, or included on a multichip package that also contains one or more processors. In some embodiments, the NIC 220 may include a local processor (not shown) and/or a local memory (not shown) that are both local to the NIC 220. In such embodiments, the local processor of the NIC 220 may be capable of performing one or more of the functions of the compute engine 210 described herein. Additionally or alternatively, in such embodiments, the local memory of the NIC 220 may be integrated into one or more components of the compute device 120 at the board level, socket level, chip level, and/or other levels.

[0022] Each data storage device 222, may be embodied as any type of device configured for short-term or long-term storage of data such as, for example, memory devices and circuits, memory cards, hard disk drives, solid-state drives, or other data storage device. Each data storage device 222 may include a system partition that stores data and firmware code for the data storage device 222 and one or more operating system partitions that store data files and executables for operating systems.

[0023] The compute device 150 may have components similar to those described in FIG. 2 with reference to the compute device 120. The description of those components of the compute

device 120 is equally applicable to the description of components of the compute device 150. Further, it should be appreciated that any of the compute devices 120, 150 may include other components, sub-components, and devices commonly found in a computing device, which are not discussed above in reference to the compute device 120 and not discussed herein for clarity of the description.

[0024] As described above, the compute devices 120, 150 are illustratively in communication via the network 112, which may be embodied as any type of wired or wireless communication network, including global networks (e.g., the Internet), local area networks (LANs) or wide area networks (WANs), cellular networks (e.g., Global System for Mobile Communications (GSM), 3G, Long Term Evolution (LTE), Worldwide Interoperability for Microwave Access (WiMAX), etc.), a radio area network (RAN), digital subscriber line (DSL) networks, cable networks (e.g., coaxial networks, fiber networks, etc.), or any combination thereof.

[0025] Referring now to FIG. 3, the compute device 120, in operation, may execute a method 300 for fusing data from multiple sensors. The method 300 is also illustrated as a pipeline 600 in FIG. 6. The method 300 begins with block 302, in which the compute device 120 determines whether to enable sensor fusion. In making the determination, the compute device 120 may determine whether a configuration setting (e.g., stored in a configuration file in the data storage device 222) indicates that sensor fusion should be enabled, whether the sensor fusion logic unit 130 is operable and in communication with the sensors 140, and/or based on other factors. Regardless, in response to a determination to enable sensor fusion, the method 300 advances to block 304, in which the compute device 120 obtains detection data from multiple sensors (e.g., the sensors 140). The detection data may be embodied as any data provided by a corresponding sensor 140 and indicative of the presence of one or more objects in an environment of the host system 110. In doing so, the compute device 120 may obtain camera detection data (e.g., from the camera 142), as indicated in block 306. In the illustrative embodiment, in obtaining the camera detection data, the compute device 120 obtains two or three dimensional image data of one or more visible objects, as indicated in block 308. Additionally, in the illustrative embodiment, the compute device 120 obtains lidar detection data (e.g., from the lidar sensor 144), as indicated in block 310. In doing so, in the illustrative embodiment, the compute device 120 obtains three dimensional data that indicates the depth (e.g., in addition to x and y coordinates) of each of one or more objects in an environment of the host system 110, as indicated in block 312. In the illustrative embodiment, the lidar detection data may be embodied as a point cloud (e.g., a set of points in three dimensional space representing positions in the environment where laser pulses were reflected back to the lidar sensor 144).

[0026] As indicated in block 314, the compute device 120 may additionally obtain detection data from one or more other sensors 146. For example, the compute device 120 may obtain detection data from a radar sensor, as indicated in block 316 and/or from an acoustic sensor, as indicated in block 318. Regardless, in the illustrative embodiment, the compute device 120 synchronizes the detection data from the multiple sensors 140, as indicated in block 320. In doing so, in the illustrative embodiment, the compute device 120 synchronizes the detection data received from each sensor 140 based on time data indicative of when each set of detection data was obtained (e.g., by each corresponding sensor 140), as indicated in block 322. For example, the detection data from each sensor 140 may include a time stamp (e.g., any data indicative of the time when the detection data was produced by the corresponding sensor 140) and the compute device 120 may associate detection data from each sensor 140 with a corresponding moment in time (e.g., a given time, plus or minus a predefined amount of time), referred to herein as a frame. The compute device 120, in the illustrative embodiment, also removes noise from the obtained detection data (e.g., on a sensor-by-sensor basis), such as by performing feature extraction, as indicated in blocks 324 and 326. Further, the compute device 120, in the illustrative embodiment, performs object detection using the extracted features for each sensor 140, as indicated in block 328, and defines bounding boxes (e.g., rectangles that enclose the detected objects) for the detected objects in the detection data, as indicated in block 330. Subsequently, the method 300 advances to block 332 of FIG. 4, in which the compute device 120 merges (e.g., fuses) the detection data from the multiple sensors 140 to define final bounding shapes.

[0027] Referring now to FIG. 4, in merging the detection data from the multiple sensors 140, the compute device 120, in the illustrative embodiment, applies weights (e.g., numerical factors that selectively increase or decrease the significance of the detection data from a given sensor 140 to the total set of detection data from the sensors 140), as indicated in block 334. In doing so, and as indicated in block 336, the compute device 120 may apply the weights as a function of a range of a detected object from the corresponding sensor 140. For example, and as indicated in the chart 700 of FIG. 7, in which an 'H' represents a relatively heavy weight being applied to the detection data from the corresponding sensor 140, if the object is relatively near the host system 110 (e.g., less than a predefined distance), the detection data from the lidar sensor 144 may be more heavily weighted relative to sensor data from the camera 142 or other sensors 146. Conversely, if the object is relatively far away (e.g., greater than or equal to the predefined distance), the compute device 120 may apply a heavier weight to the detection data from the camera 142 than to the detection data from the lidar sensor 144. As indicated in block 338, the compute device 120 may apply weights as a function of the present weather in the environment of the host system 110. For example, in rain, fog, or snow, the compute device 120 may apply

heavier weights to detection data from a radar sensor or acoustic sensor than to detection data from the camera 142 or the lidar sensor 144. As another example, and as indicated in block 340, the compute device 120 may apply weights as a function of the amount of light present in the environment. In doing so, the compute device 120 may determine the amount of weight to apply to the detection data from a given sensor 140 as a function of the time of day when the detection data was obtained, as indicated in block 342. As shown in the chart 700, the compute device 120 may apply a greater (e.g., heavier) weight the detection data from the camera 142 in the day time and lesser weight to the detection data from the camera 142 in the night time. As indicated in block 344, the compute device 120 may adjust the amount of weight applied to the detection data from the camera in a direct relationship with the amount of light present (e.g., as the amount of light increases, the weight applied to the detection data from the camera increases correspondingly). As indicated in block 346, in addition to selectively applying different weights to the detection data from the different sensors 140, the compute device 120 may also apply user-defined rules (e.g., stored in a data storage device 222 or present in the memory 214) that define how to combine bounding boxes (e.g., the bounding boxes from block 330 of FIG. 3) associated with the detection data from the various sensors 140. As indicated in block 348, the compute device 120 may determine a class of each detected object (e.g., a pedestrian, a vehicle, etc.) and, as indicated in block 350, may apply one or more rules for combining bounding boxes that are specific to a predefined class of object (e.g., a specific rule for combining bounding boxes associated with vehicles).

[0028] Referring now to FIG. 5, in merging the detection data, the compute device 120 may apply a belief function that produces a degree of belief that multiple bounding boxes (e.g., bounding boxes present in the detection data from multiple sensors 140) refer to the same object, different objects, or free space, as indicated in block 352. In doing so, the compute device 120 may determine a degree of belief as a function of an area of overlap between multiple bounding boxes, as indicated in block 354. For example, the compute device may determine a degree of belief as a function of an intersection-over-union (IOU) between multiple bounding boxes, as indicated in block 356. The compute device 120 may determine the IOU by calculating the area of intersection between the bounding boxes, divided by the area of the union of the bounding boxes. In the illustrative embodiment, the compute device 120 determines a degree of belief that multiple bounding boxes represent the same object, a related degree of belief that the bounding boxes represent different objects, and a degree of belief that the bounding boxes are associated with free space, represented by a false positive rate. Table 1, below, illustrates an embodiment of the calculations made by the compute device 120 to determine each corresponding degree of belief:

Hypothesis	Degree of Belief
Same Object	$(1 - FalsePositiveRate) * iou(box_1, box_2)$
Different Object	$(1 - FalsePositiveRate) * (1 - iou(box_1, box_2))$
Free Space	$FalsePositiveRate$

Table 1. IOU Degrees of Belief Calculations.

[0029] The compute device 120 may additionally or alternatively determine a degree of belief as a function of an intersection-over-minimum (IOM) between multiple bounding boxes, as indicated in block 358. For example, in the illustrative embodiment, the compute device 120 may determine one or more degrees of belief based on an IOM if the object associated with the bounding boxes has been determined to be a vehicle, as the IOM is particularly useful in correcting lidar detection data, which may include multiple relatively small bounding boxes for a single vehicle (e.g., as a result of the lidar's laser pulses penetrating the windows of the vehicle). Table 2, below, illustrates an embodiment of the calculations made by the compute device 120 to determine each degree of belief (e.g., same object, different object, and free space):

Hypothesis	Degree of Belief
Same Object	$(1 - FalsePositiveRate) * iom(box_1, box_2)$
Different Object	$(1 - FalsePositiveRate) * (1 - iom(box_1, box_2))$
Free Space	$FalsePositiveRate$

Table 2. IOM Degrees of Belief Calculations.

[0030] As indicated in block 360, the compute device 120 may determine a degree of belief as a function of a proximity of objects in multiple bounding boxes to each other. Determining the degree of belief based on proximity of objects to each other is particularly useful in detecting distinct objects in crowded scenes. In doing so, the compute device 120 may determine the proximity of the objects based on depth data from the lidar detection data or from the 3D camera detection data (e.g., by determining differences in the depths of different bounding boxes), as indicated in block 362. Alternatively, the compute device 120 may determine distances between objects based on x-y coordinate pairs representing the center of each bounding box, as indicated in block 364. An example function, given in Equation 1 below, for calculating a degree of belief (a "proximity mass") using proximity may utilize a depth threshold (T) that specifies a minimum separation between two distinct objects, where depth is the distance from the host system 110 as measured by the lidar sensor 144.

[0031] $$ProximityMass = \max\left(\frac{abs(depth_1, depth_2)}{T}, 1.0\right) \quad \text{(Equation 1)}$$

[0032] Corresponding degrees of belief for a hypothesis that multiple bounding boxes represent the same object, a hypothesis that the multiple bounding boxes represent different objects, or that the bounding boxes represent free space are given below in Table 3. The result may also be unknown, as indicated below.

Hypothesis	**Degree of Belief**
Same Object	$(1 - FalsePositiveRate) * 1 - ProximityMass$
Different Object	$(1 - FalsePositiveRate) * ProximityMass$
Free Space	$FalsePositiveRate$
Unknown	*1.0 if ($depth_1$ or $depth_2$ is undefined)*

Table 3. Proximity Degrees of Belief Calculations.

[0033] In the illustrative embodiment, the compute device 120 calculates degrees of belief using the Dempster-Shafer rule (e.g., calculating a degree of belief for different possible hypotheses) and determines which hypothesis has the greatest degree of belief (e.g., that the bounding boxes represent the same object, that the bounding boxes represent different objects, etc.), as indicated in block 366. As indicated in block 368, the compute device 120 may apply an object localization policy that redefines an object area when the detection data for different sensors indicate different boundaries for the same object. In the illustrative embodiment, and as indicated in block 370, the compute device 120 may perform a search algorithm (e.g., an O(N) algorithm, in which N represents a number of bounding boxes in the detection data and k represents the neighborhood or localization distance for the search) to determine whether to fuse (e.g., merge) bounding boxes. In doing so, the compute device 120 may perform an initialization operation in which the compute device 120 sets a reconciled combined detections list (d) to an empty list. Further, the compute device 120 may sort the bounding boxes from the detection data of the multiple sensors (e.g., sorting the bounding boxes by their x-axis values). Additionally, the compute device 120 may perform a set of operations for each bounding box, s, in the list of sorted boxes, as follows: (i) compute degrees of belief (e.g., masses) using different sensor fusion policies for each bounding box, s, and bounding boxes in combined detections, t, that are within a localization distance (e.g., a predefined distance representing the neighborhood of the objects); (ii) combine masses using the Dempster-Shafer rule and compute the maximum belief; (iii) if the maximum belief is that the bounding boxes refer to the same object, then fuse (e.g., merge) the bounding boxes, s and t. Otherwise, add bounding box, s, as a distinct object in the list of reconciled combined detections, d. In defining a final bounding shape (e.g., from fusing multiple bounding boxes together), the compute device 120 may define the final bounding shape as a box or other polygon (e.g., a non-rectangular polygon), as indicated in block 372. In some embodiments, the compute device 120 may utilize sequences of frames (e.g., detection data from

each of multiple sequential moments in time) to increase a confidence in the determination of the final bounding shapes, as indicated in block 374. Afterwards, the method 300 loops back to block 302, in which the compute device 120 determines whether to continue to perform sensor fusion. It should be understood that the compute device 120 may perform other operations, including determining responsive actions that the host system 110 should take based on an identification of a particular object in the environment, which are not shown in the method 300. Further, it should be understood that while the method 300 has been shown in a particular sequence, the compute device 120 may perform operations in the method 300 in a different order and/or concurrently (e.g., in parallel). Additionally, and as discussed with reference to FIG. 1, a remote compute device (e.g., the compute device 150) in communication with the host system 110 (e.g., through the network 112) may perform one or more of the operations of the method 300 instead of, or in cooperation with, the compute device 120.

EXAMPLES

[0034] Illustrative examples of the technologies disclosed herein are provided below. An embodiment of the technologies may include any one or more, and any combination of, the examples described below.

[0035] Example 1 includes a compute device comprising circuitry to obtain detection data indicative of objects detected by one or more sensors of a host system, wherein the detection data includes camera detection data indicative of at least one of a two or three dimensional image of detected objects and lidar detection data indicative of depths of detected objects; and merge the detection data from the multiple sensors to define final bounding shapes for the detected objects.

[0036] Example 2 includes the subject matter of Example 1, and wherein to merge the detection data from the one or more sensors comprises to apply weights to the detection data as a function of a context in which the detection data was obtained.

[0037] Example 3 includes the subject matter of any of Examples 1 and 2, and wherein to apply weights to the detection data comprises to apply weights as a function of a range of a detected object from the corresponding sensor.

[0038] Example 4 includes the subject matter of any of Examples 1-3, and wherein to apply weights to the detection data comprises to apply weights as a function of a present weather in an environment of the host system.

[0039] Example 5 includes the subject matter of any of Examples 1-4, and wherein to apply weights to the detection data comprises to apply weights as a function of an amount of light present in an environment of the host system.

[0040] Example 6 includes the subject matter of any of Examples 1-5, and wherein to apply weights to the detection data comprises to adjust a weight applied to the camera detection data in a direct relationship with the amount of light.

[0041] Example 7 includes the subject matter of any of Examples 1-6, and wherein to merge the detection data from the one or more sensors comprises to apply one or more user-defined rules to combine bounding boxes present in the detection data from each sensor.

[0042] Example 8 includes the subject matter of any of Examples 1-7, and wherein to merge the detection data from the one or more sensors comprises to determine a class of an object in the detection data; and apply a rule defined to combine bounding boxes for objects of the determined class.

[0043] Example 9 includes the subject matter of any of Examples 1-8, and wherein to merge the detection data from the one or more sensors comprises to apply a belief function that produces a degree of belief that multiple bounding boxes refer to the same object, different objects, or free space.

[0044] Example 10 includes the subject matter of any of Examples 1-9, and wherein to determine the degree of belief comprises to determine the degree of belief as a function of an intersection-over-union between multiple bounding boxes, determine the degree of belief as a function of an intersection-over-minimum between multiple bounding boxes, or determine a degree of belief as a function of a proximity of the objects to each other.

[0045] Example 11 includes the subject matter of any of Examples 1-10, and wherein to merge the detection data further comprises to determine whether the degree of belief that the multiple bounding boxes refer to the same object is greater than the degree of belief that the bounding boxes refer to the different objects and the degree of belief that the bounding boxes refer to free space, and merge, in response to a determination that the greatest belief is that the bounding boxes refer to the same object, the bounding boxes.

[0046] Example 12 includes the subject matter of any of Examples 1-11, and wherein to merge the bounding boxes comprises to define the final bounding shape as a polygon.

[0047] Example 13 includes one or more machine-readable storage media comprising a plurality of instructions stored thereon that, in response to being executed, cause a compute device to obtain detection data indicative of objects detected by each of one or more sensors of a host system, wherein the detection data includes camera detection data indicative of at least one of a two or three dimensional image of detected objects and lidar detection data indicative of depths of detected objects; and merge the detection data from the multiple sensors to define final bounding shapes for the detected objects.

[0048] Example 14 includes the subject matter of Example 13, and wherein to merge the detection data from the one or more sensors comprises to apply weights to the detection data as a function of a context in which the detection data was obtained.

[0049] Example 15 includes the subject matter of any of Examples 13 and 14, and wherein to apply weights to the detection data comprises to apply weights as a function of a range of a detected object from the corresponding sensor.

[0050] Example 16 includes the subject matter of any of Examples 13-15, and wherein to apply weights to the detection data comprises to apply weights as a function of a present weather in an environment of the host system.

[0051] Example 17 includes the subject matter of any of Examples 13-16, and wherein to apply weights to the detection data comprises to apply weights as a function of an amount of light present in an environment of the host system.

[0052] Example 18 includes the subject matter of any of Examples 13-17, and wherein to merge the detection data from the one or more sensors comprises to determine a class of an object in the detection data; and apply a rule defined to combine bounding boxes for objects of the determined class.

[0053] Example 19 includes a method comprising obtaining, by a compute device, detection data indicative of objects detected by one or more sensors of a host system, wherein the detection data includes camera detection data indicative of at least one of a two or three dimensional image of detected objects and lidar detection data indicative of depths of detected objects; and merging, by the compute device, the detection data from the multiple sensors to define final bounding shapes for the detected objects.

[0054] Example 20 includes the subject matter of Example 19, and wherein merging the detection data from the one or more sensors comprises applying weights to the detection data as a function of a context in which the detection data was obtained.

WHAT IS CLAIMED IS:

1. A compute device comprising:

circuitry to:

obtain detection data indicative of objects detected by one or more sensors of a host system, wherein the detection data includes camera detection data indicative of at least one of a two or three dimensional image of detected objects and lidar detection data indicative of depths of detected objects; and

merge the detection data from the multiple sensors to define final bounding shapes for the detected objects.

2. The compute device of claim 1, wherein to merge the detection data from the one or more sensors comprises to apply weights to the detection data as a function of a context in which the detection data was obtained.

3. The compute device of claim 2, wherein to apply weights to the detection data comprises to apply weights as a function of a range of a detected object from the corresponding sensor.

4. The compute device of claim 2, wherein to apply weights to the detection data comprises to apply weights as a function of a present weather in an environment of the host system.

5. The compute device of claim 2, wherein to apply weights to the detection data comprises to apply weights as a function of an amount of light present in an environment of the host system.

6. The compute device of claim 5, wherein to apply weights to the detection data comprises to adjust a weight applied to the camera detection data in a direct relationship with the amount of light.

7. The compute device of claim 1, wherein to merge the detection data from the one or more sensors comprises to apply one or more user-defined rules to combine bounding boxes present in the detection data from each sensor.

8. The compute device of claim 1, wherein to merge the detection data from the one or more sensors comprises to:

determine a class of an object in the detection data; and

apply a rule defined to combine bounding boxes for objects of the determined class.

9. The compute device of claim 1, wherein to merge the detection data from the one or more sensors comprises to apply a belief function that produces a degree of belief that multiple bounding boxes refer to the same object, different objects, or free space.

10. The compute device of claim 9, wherein to determine the degree of belief comprises to determine the degree of belief as a function of an intersection-over-union between multiple bounding boxes, determine the degree of belief as a function of an intersection-over-minimum between multiple bounding boxes, or determine a degree of belief as a function of a proximity of the objects to each other.

11. The compute device of claim 10, wherein to merge the detection data further comprises to determine whether the degree of belief that the multiple bounding boxes refer to the same object is greater than the degree of belief that the bounding boxes refer to the different objects and the degree of belief that the bounding boxes refer to free space, and merge, in response to a determination that the greatest belief is that the bounding boxes refer to the same object, the bounding boxes.

12. The compute device of claim 11, wherein to merge the bounding boxes comprises to define the final bounding shape as a polygon.

13. A method comprising:

obtaining, by a compute device, detection data indicative of objects detected by one or more sensors of a host system, wherein the detection data includes camera detection data indicative of at least one of a two or three dimensional image of detected objects and lidar detection data indicative of depths of detected objects; and

merging, by the compute device, the detection data from the multiple sensors to define final bounding shapes for the detected objects.

14. The method of claim 13, wherein merging the detection data from the one or more sensors comprises applying weights to the detection data as a function of a context in which the detection data was obtained.

15. The method of claim 14, wherein applying weights to the detection data comprises applying weights as a function of a range of a detected object from the corresponding sensor.

16. The method of claim 14, wherein applying weights to the detection data comprises applying weights as a function of a present weather in an environment of the host system.

17. The method of claim 14, wherein applying weights to the detection data comprises applying weights as a function of an amount of light present in an environment of the host system.

18. The method of claim 17, wherein applying weights to the detection data comprises adjusting a weight applied to the camera detection data in a direct relationship with the amount of light.

19. The method of claim 13, wherein merging the detection data from the one or more sensors comprises applying one or more user-defined rules to combine bounding boxes present in the detection data from each sensor.

20. The method of claim 13, wherein merging the detection data from the one or more sensors comprises to:

determining a class of an object in the detection data; and

applying a rule defined to combine bounding boxes for objects of the determined class.

21. The method of claim 13, wherein merging the detection data from the one or more sensors comprises applying a belief function that produces a degree of belief that multiple bounding boxes refer to the same object, different objects, or free space.

22. The method of claim 21, wherein determining the degree of belief comprises to determine the degree of belief as a function of an intersection-over-union between multiple bounding boxes, determining the degree of belief as a function of an intersection-over-minimum between multiple bounding boxes, or determining a degree of belief as a function of a proximity of the objects to each other.

23. The method of claim 22, wherein merging the detection data further comprises determining whether the degree of belief that the multiple bounding boxes refer to the same object is greater than the degree of belief that the bounding boxes refer to the different objects and the degree of belief that the bounding boxes refer to free space, and merging, in response to a determination that the greatest belief is that the bounding boxes refer to the same object, the bounding boxes.

24. The method of claim 23, wherein merging the bounding boxes comprises defining the final bounding shape as a polygon.

25. One or more machine-readable storage media comprising a plurality of instructions stored thereon that, in response to being executed, cause a compute device to perform the method of any of claims 13-24.

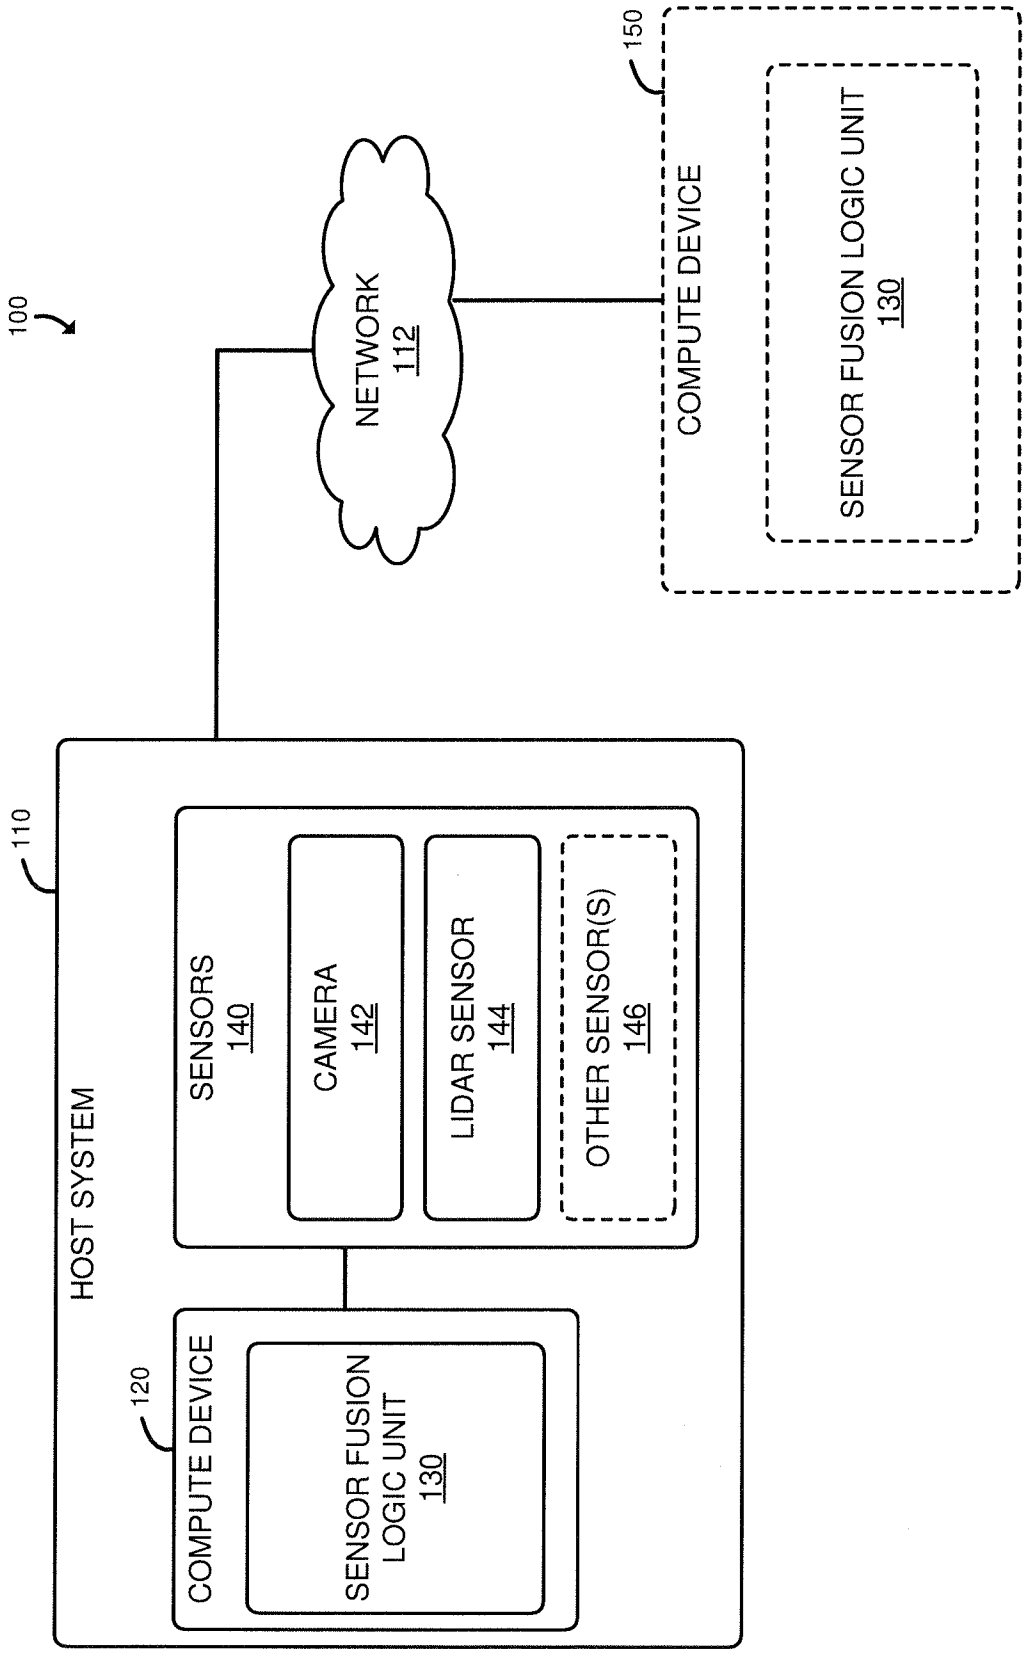
100
NETWORK
112
150
COMPUTE DEVICE
SENSOR FUSION LOGIC UNIT
130
110
HOST SYSTEM
SENSORS
140
CAMERA
142
LIDAR SENSOR
144
OTHER SENSOR(S)
146
120
COMPUTE DEVICE
SENSOR FUSION LOGIC UNIT
130

FIG. 1

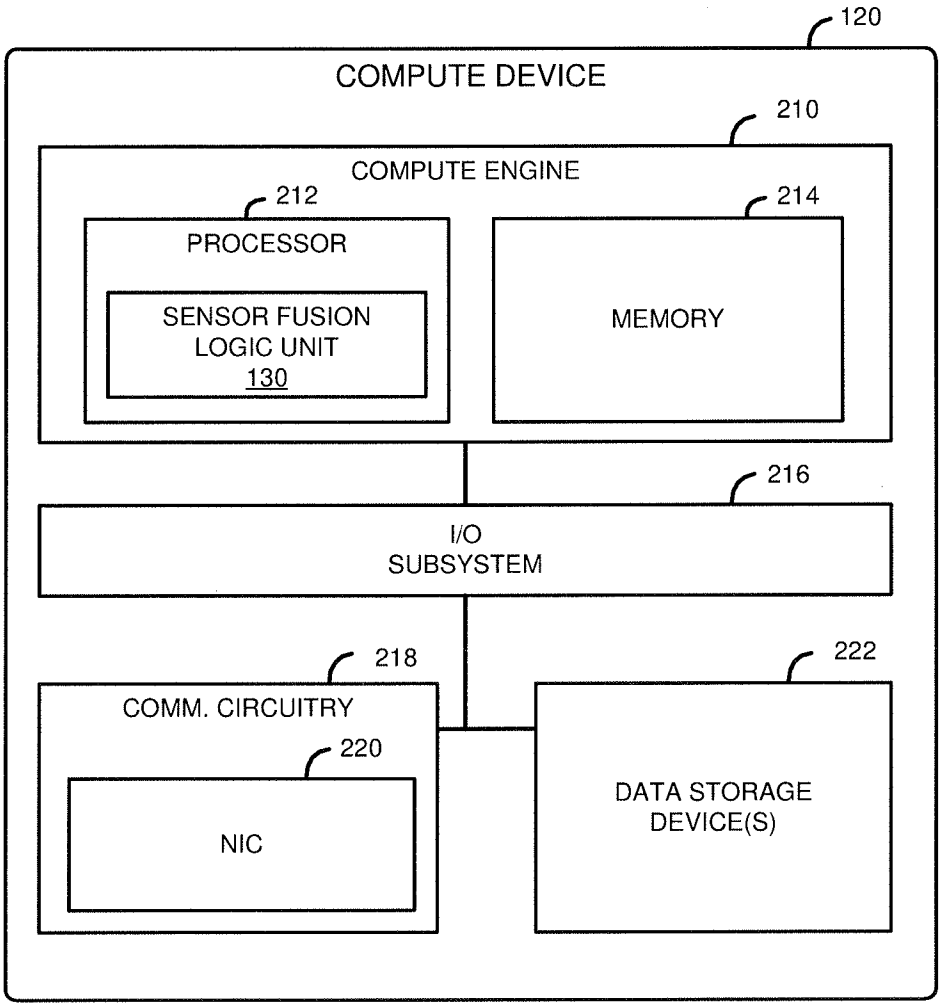
120
COMPUTE DEVICE
210
COMPUTE ENGINE
212
PROCESSOR
SENSOR FUSION
LOGIC UNIT
130
214
MEMORY
216
I/O
SUBSYSTEM
218
COMM. CIRCUITRY
220
NIC
222
DATA STORAGE
DEVICE(S)

FIG. 2

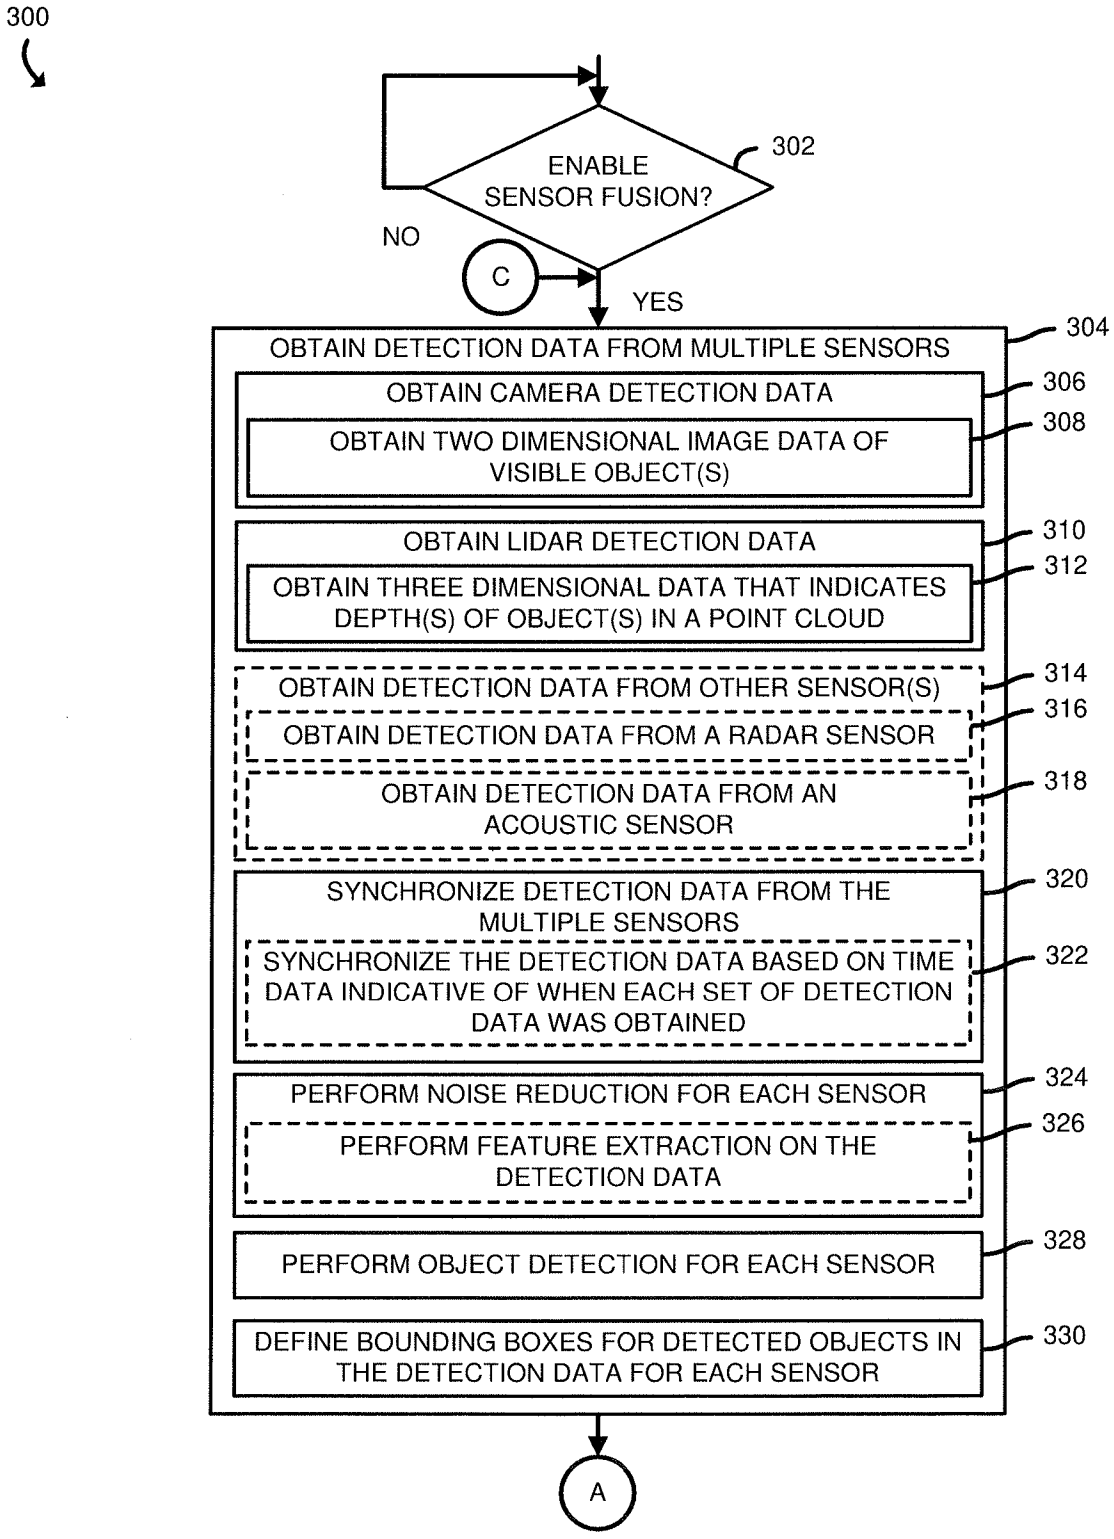
300
ENABLE SENSOR FUSION?
302
NO
C
YES
OBTAIN DETECTION DATA FROM MULTIPLE SENSORS
304
OBTAIN CAMERA DETECTION DATA
306
OBTAIN TWO DIMENSIONAL IMAGE DATA OF VISIBLE OBJECT(S)
308
OBTAIN LIDAR DETECTION DATA
310
OBTAIN THREE DIMENSIONAL DATA THAT INDICATES DEPTH(S) OF OBJECT(S) IN A POINT CLOUD
312
OBTAIN DETECTION DATA FROM OTHER SENSOR(S)
314
OBTAIN DETECTION DATA FROM A RADAR SENSOR
316
OBTAIN DETECTION DATA FROM AN ACOUSTIC SENSOR
318
SYNCHRONIZE DETECTION DATA FROM THE MULTIPLE SENSORS
320
SYNCHRONIZE THE DETECTION DATA BASED ON TIME DATA INDICATIVE OF WHEN EACH SET OF DETECTION DATA WAS OBTAINED
322
PERFORM NOISE REDUCTION FOR EACH SENSOR
324
PERFORM FEATURE EXTRACTION ON THE DETECTION DATA
326
PERFORM OBJECT DETECTION FOR EACH SENSOR
328
DEFINE BOUNDING BOXES FOR DETECTED OBJECTS IN THE DETECTION DATA FOR EACH SENSOR
330
A

FIG. 3

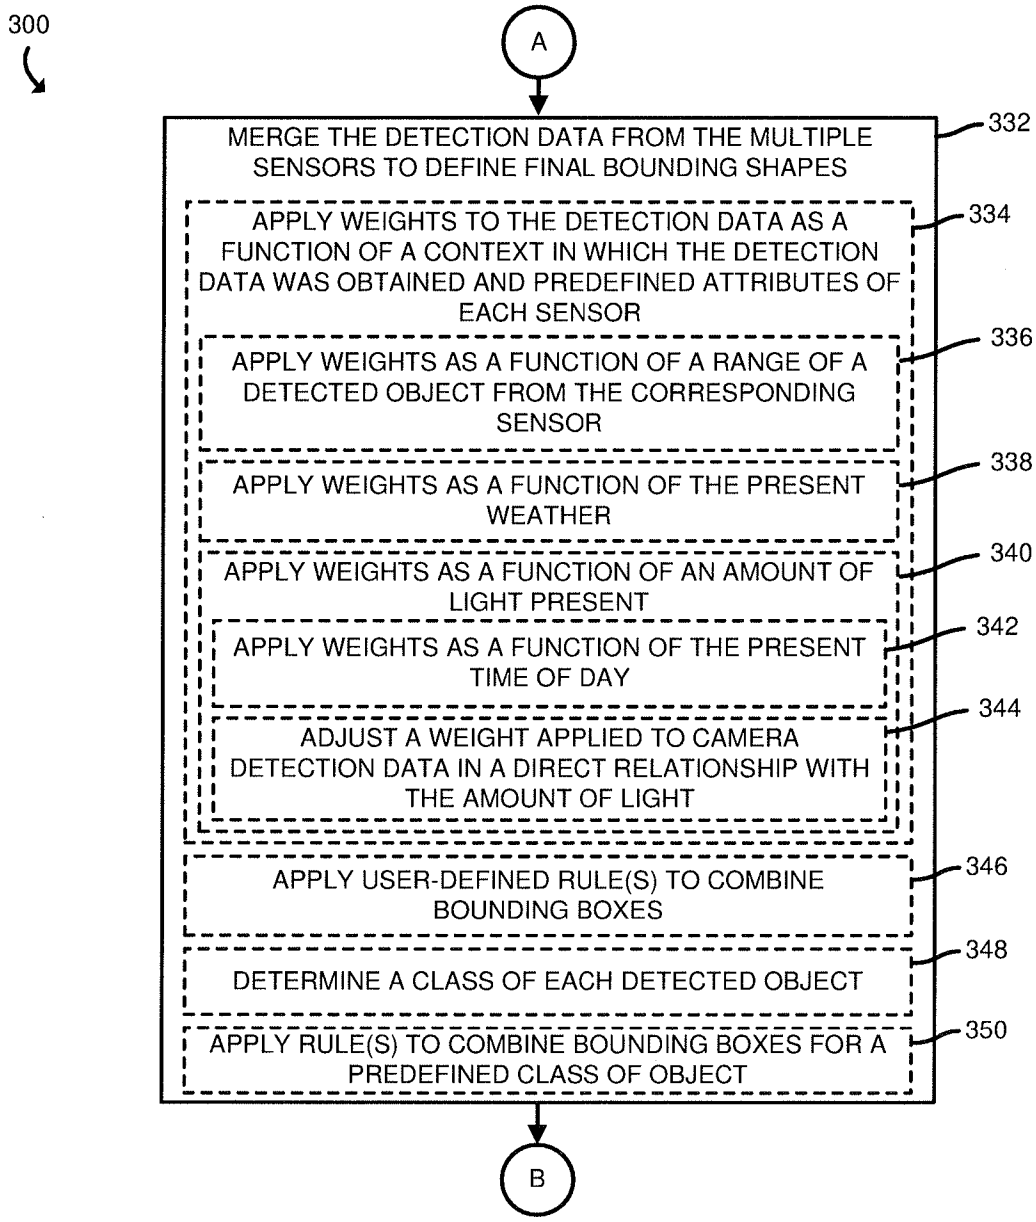
300
A
MERGE THE DETECTION DATA FROM THE MULTIPLE SENSORS TO DEFINE FINAL BOUNDING SHAPES
332
APPLY WEIGHTS TO THE DETECTION DATA AS A FUNCTION OF A CONTEXT IN WHICH THE DETECTION DATA WAS OBTAINED AND PREDEFINED ATTRIBUTES OF EACH SENSOR
334
APPLY WEIGHTS AS A FUNCTION OF A RANGE OF A DETECTED OBJECT FROM THE CORRESPONDING SENSOR
336
APPLY WEIGHTS AS A FUNCTION OF THE PRESENT WEATHER
338
APPLY WEIGHTS AS A FUNCTION OF AN AMOUNT OF LIGHT PRESENT
340
APPLY WEIGHTS AS A FUNCTION OF THE PRESENT TIME OF DAY
342
ADJUST A WEIGHT APPLIED TO CAMERA DETECTION DATA IN A DIRECT RELATIONSHIP WITH THE AMOUNT OF LIGHT
344
APPLY USER-DEFINED RULE(S) TO COMBINE BOUNDING BOXES
346
DETERMINE A CLASS OF EACH DETECTED OBJECT
348
APPLY RULE(S) TO COMBINE BOUNDING BOXES FOR A PREDEFINED CLASS OF OBJECT
350
B

FIG. 4

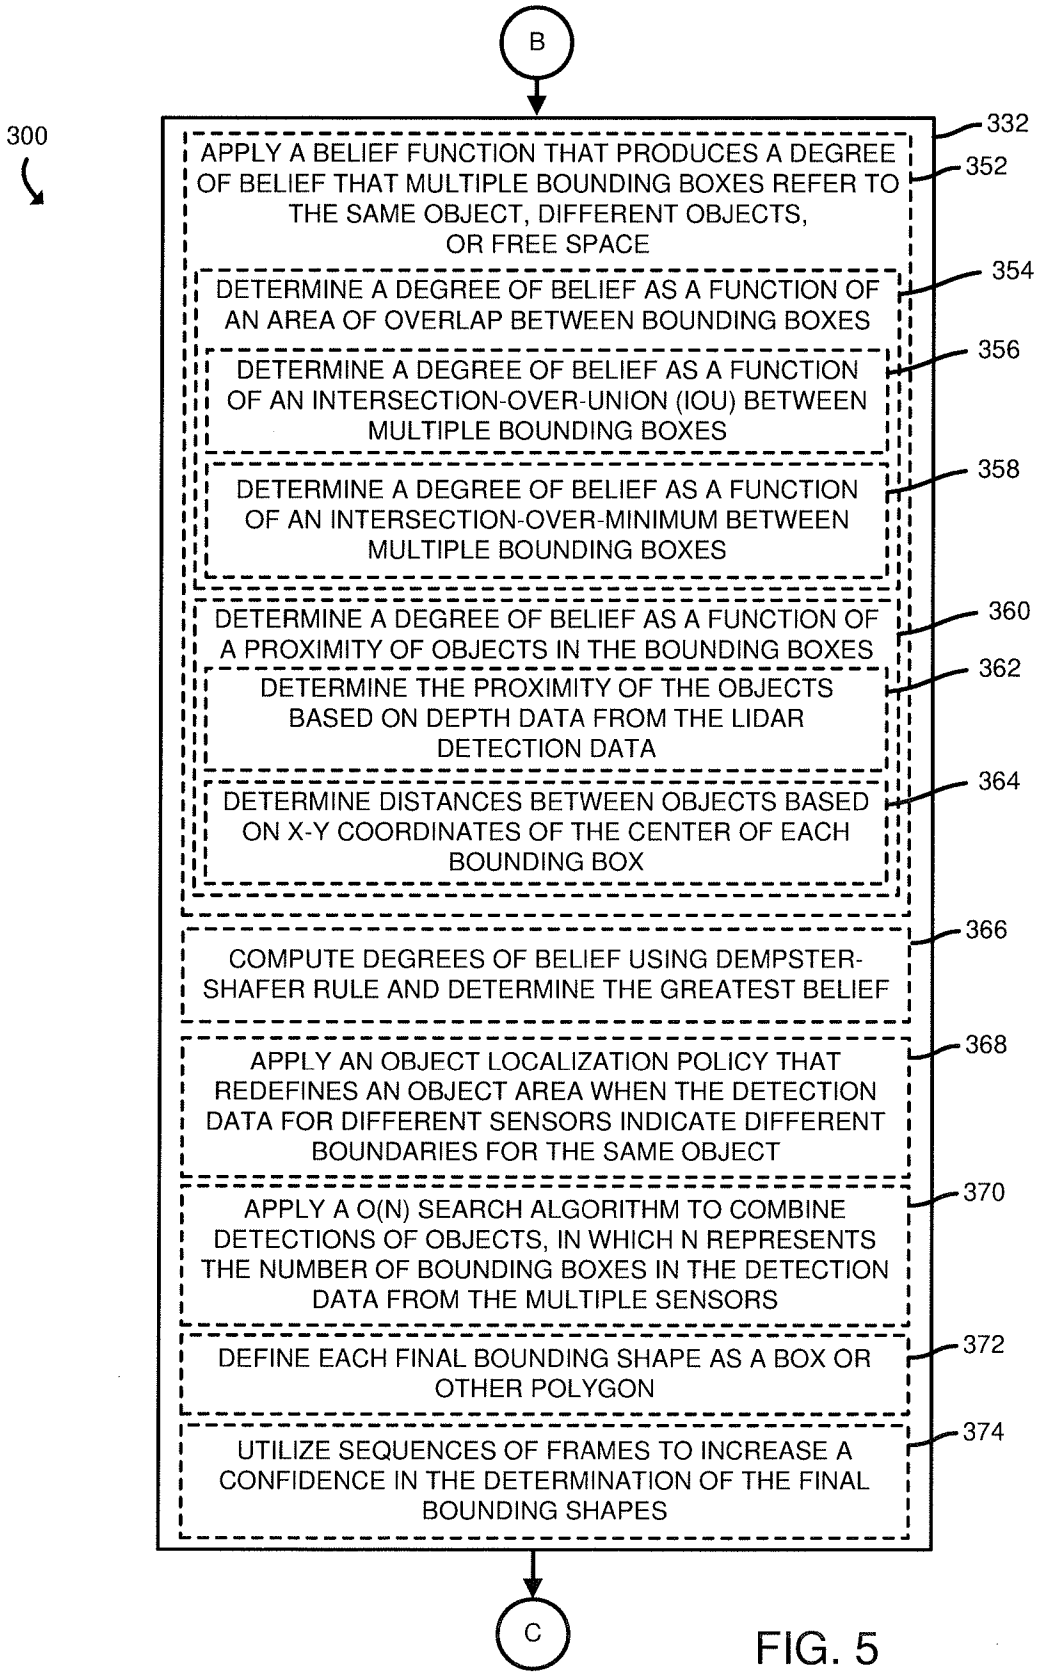
B
300
APPLY A BELIEF FUNCTION THAT PRODUCES A DEGREE OF BELIEF THAT MULTIPLE BOUNDING BOXES REFER TO THE SAME OBJECT, DIFFERENT OBJECTS, OR FREE SPACE
332
352
DETERMINE A DEGREE OF BELIEF AS A FUNCTION OF AN AREA OF OVERLAP BETWEEN BOUNDING BOXES
354
DETERMINE A DEGREE OF BELIEF AS A FUNCTION OF AN INTERSECTION-OVER-UNION (IOU) BETWEEN MULTIPLE BOUNDING BOXES
356
DETERMINE A DEGREE OF BELIEF AS A FUNCTION OF AN INTERSECTION-OVER-MINIMUM BETWEEN MULTIPLE BOUNDING BOXES
358
DETERMINE A DEGREE OF BELIEF AS A FUNCTION OF A PROXIMITY OF OBJECTS IN THE BOUNDING BOXES
360
DETERMINE THE PROXIMITY OF THE OBJECTS BASED ON DEPTH DATA FROM THE LIDAR DETECTION DATA
362
DETERMINE DISTANCES BETWEEN OBJECTS BASED ON X-Y COORDINATES OF THE CENTER OF EACH BOUNDING BOX
364
COMPUTE DEGREES OF BELIEF USING DEMPSTER-SHAFER RULE AND DETERMINE THE GREATEST BELIEF
366
APPLY AN OBJECT LOCALIZATION POLICY THAT REDEFINES AN OBJECT AREA WHEN THE DETECTION DATA FOR DIFFERENT SENSORS INDICATE DIFFERENT BOUNDARIES FOR THE SAME OBJECT
368
APPLY A O(N) SEARCH ALGORITHM TO COMBINE DETECTIONS OF OBJECTS, IN WHICH N REPRESENTS THE NUMBER OF BOUNDING BOXES IN THE DETECTION DATA FROM THE MULTIPLE SENSORS
370
DEFINE EACH FINAL BOUNDING SHAPE AS A BOX OR OTHER POLYGON
372
UTILIZE SEQUENCES OF FRAMES TO INCREASE A CONFIDENCE IN THE DETERMINATION OF THE FINAL BOUNDING SHAPES
374
C

FIG. 5

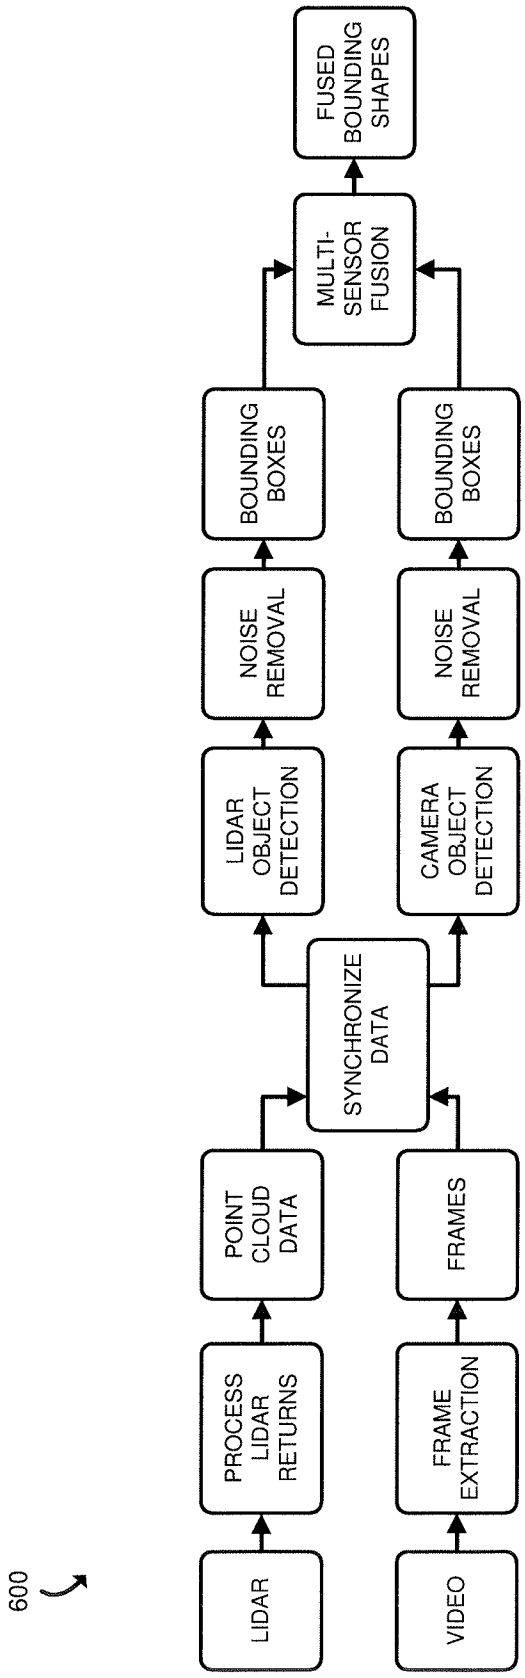
600
LIDAR
VIDEO
PROCESS LIDAR RETURNS
FRAME EXTRACTION
POINT CLOUD DATA
FRAMES
SYNCHRONIZE DATA
LIDAR OBJECT DETECTION
CAMERA OBJECT DETECTION
NOISE REMOVAL
NOISE REMOVAL
BOUNDING BOXES
BOUNDING BOXES
MULTI-SENSOR FUSION
FUSED BOUNDING SHAPES

FIG. 6

700

Relative Weights of the Sensors under Different Contexts

	CONTEXT									
Sensors	Day	Night	Nearer Range/ Lower elevation	Far ranges	police siren, ambulance, firetruck, explosion	Tunnel	Road condition (vibration sensors)	Rain/Fog/Snow	Fire hazards on the road/ Forest fire	Glare/ Shadows
LiDAR		H	H			H	H		H	H
Camera	H			H						
Radar				H				H	H	H
Acoustic					H			H		

FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2019/061955

A. CLASSIFICATION OF SUBJECT MATTER

G06T 7/73(2017.01)i

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)
G06T 7/73; G06K 9/46; G06K 9/62; G06T 11/60; G06T 7/60; H04N 7/26

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & keywords: sensor, camera, image, light detection and ranging, depth, bounding

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	BHARANIDHAR DURAISAMY et al., `Track Level Fusion of Extended Objects from Heterogeneous Sensors', 2016 19th International Conference on Information Fusion (FUSION), 05-08 July 2016, pp. 1-10 pages 1-3; and figures 1-2	1,13,25
Y		2-12,14-24
Y	PRAVEEN KUMAR et al., `Addressing uncertainty in multi-modal fusion for improved object detection in dynamic environment', Information Fusion 11, 2010, pp. 311-324 pages 313-316; and figure 1	2-6,9-12,14-18 ,21-24
Y	US 2013-0039409 A1 (PUNEET GUPTA) 14 February 2013 paragraphs [0024]-[0025]	7,19
Y	US 2018-0047193 A1 (QUALCOMM INCORPORATED) 15 February 2018 claim 1	8,20
A	KR 10-2017-0048108 A (SAMSUNG SDS CO., LTD.) 08 May 2017 claims 1, 4; and figure 12	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	Date of mailing of the international search report
09 March 2020 (09.03.2020)	**10 March 2020 (10.03.2020)**

Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea
Facsimile No. +82-42-481-8578

Authorized officer
KIM, Sung Hoon
Telephone No. +82-42-481-8710

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2019/061955

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013-0039409 A1	14/02/2013	US 8917764 B2	23/12/2014
US 2018-0047193 A1	15/02/2018	None	
KR 10-2017-0048108 A	08/05/2017	None	