US 20180314253A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2018/0314253 A1**

Mercep et al. (43) **Pub. Date: Nov. 1, 2018**

(54) **EMBEDDED AUTOMOTIVE PERCEPTION WITH MACHINE LEARNING CLASSIFICATION OF SENSOR DATA**

(71) Applicant: **Mentor Graphics Development (Deutschland) GmbH**, Wilsonville, OR (US)

(72) Inventors: **Ljubo Mercep**, Munchen (DE); **Matthias Pollach**, Munchen (DE)

(21) Appl. No.: **15/872,259**

(22) Filed: **Jan. 16, 2018**

Related U.S. Application Data

(60) Provisional application No. 62/492,416, filed on May 1, 2017.

Publication Classification

(51) **Int. Cl.**

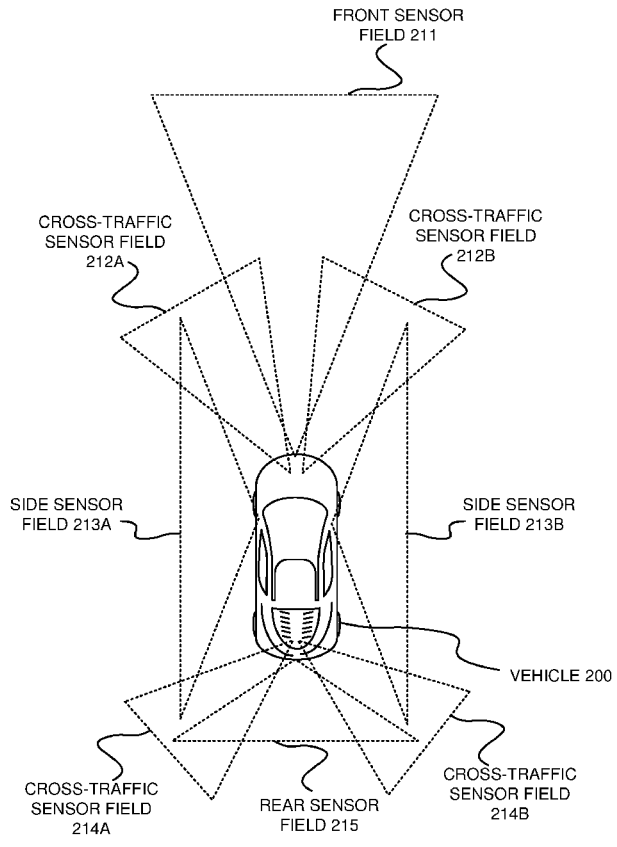
G05D 1/00	(2006.01)
G06K 9/00	(2006.01)
G06K 9/62	(2006.01)
G06K 9/66	(2006.01)
G06T 7/55	(2006.01)
G06T 7/521	(2006.01)
G06N 99/00	(2006.01)

(52) **U.S. Cl.**

CPC ***G05D 1/0088*** (2013.01); ***G06K 9/00825*** (2013.01); ***G06K 9/6296*** (2013.01); ***G06K 9/66*** (2013.01); ***G06T 7/55*** (2017.01); ***G06K 9/6289*** (2013.01); *G06T 2207/20072* (2013.01); ***G06T 7/521*** (2017.01); ***G06N 99/005*** (2013.01); *G06T 2207/10044* (2013.01); *G06T 2207/30261* (2013.01); *G06T 2207/20081* (2013.01); ***G06K 9/00805*** (2013.01)

(57) **ABSTRACT**

This application discloses a computing system to implement perception in sensor data for an assisted or automated driving system of a vehicle. The computing system can generate a matchable representation of sensor measurement data collected by sensors mounted in a vehicle, and compare the matchable representation of the sensor measurement data to an object model describing a type of an object capable of being located proximate to the vehicle. Based on the comparison, the computing system can classify the sensor measurement data as corresponding to the type of the object based, at least in part, on the comparison of the matchable representation of the sensor measurement data to the object model. A control system for the vehicle can configured to control operation of the vehicle based, at least in part, on the classified type of the object for the sensor measurement data.



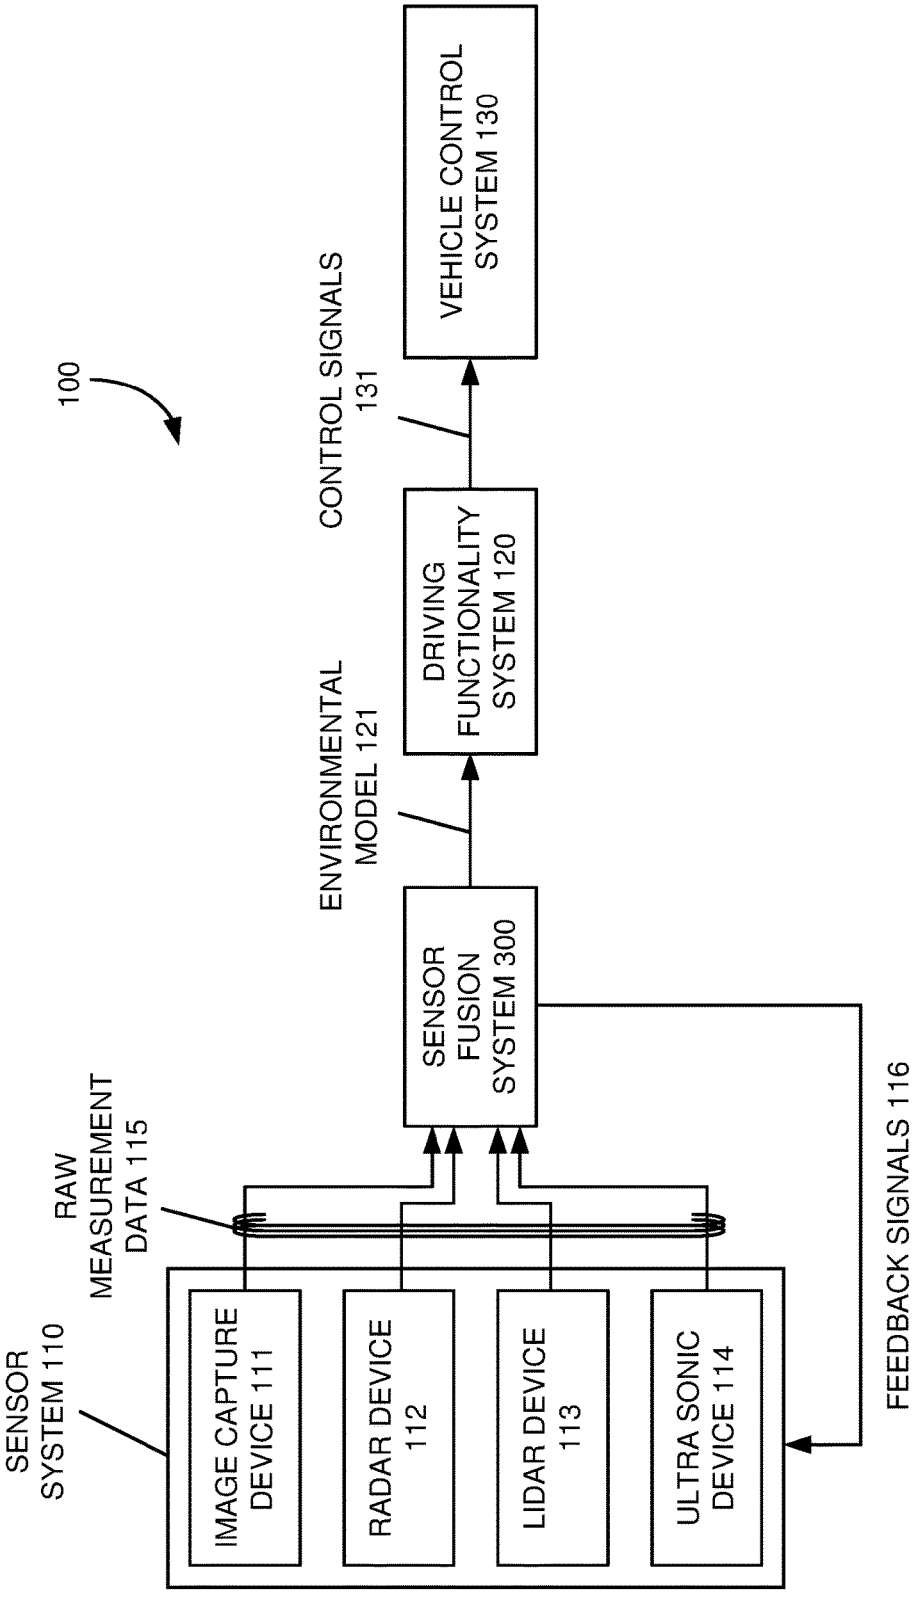

100
SENSOR SYSTEM 110
IMAGE CAPTURE DEVICE 111
RADAR DEVICE 112
LIDAR DEVICE 113
ULTRA SONIC DEVICE 114
RAW MEASUREMENT DATA 115
SENSOR FUSION SYSTEM 300
ENVIRONMENTAL MODEL 121
DRIVING FUNCTIONALITY SYSTEM 120
CONTROL SIGNALS 131
VEHICLE CONTROL SYSTEM 130
FEEDBACK SIGNALS 116

FIG. 1

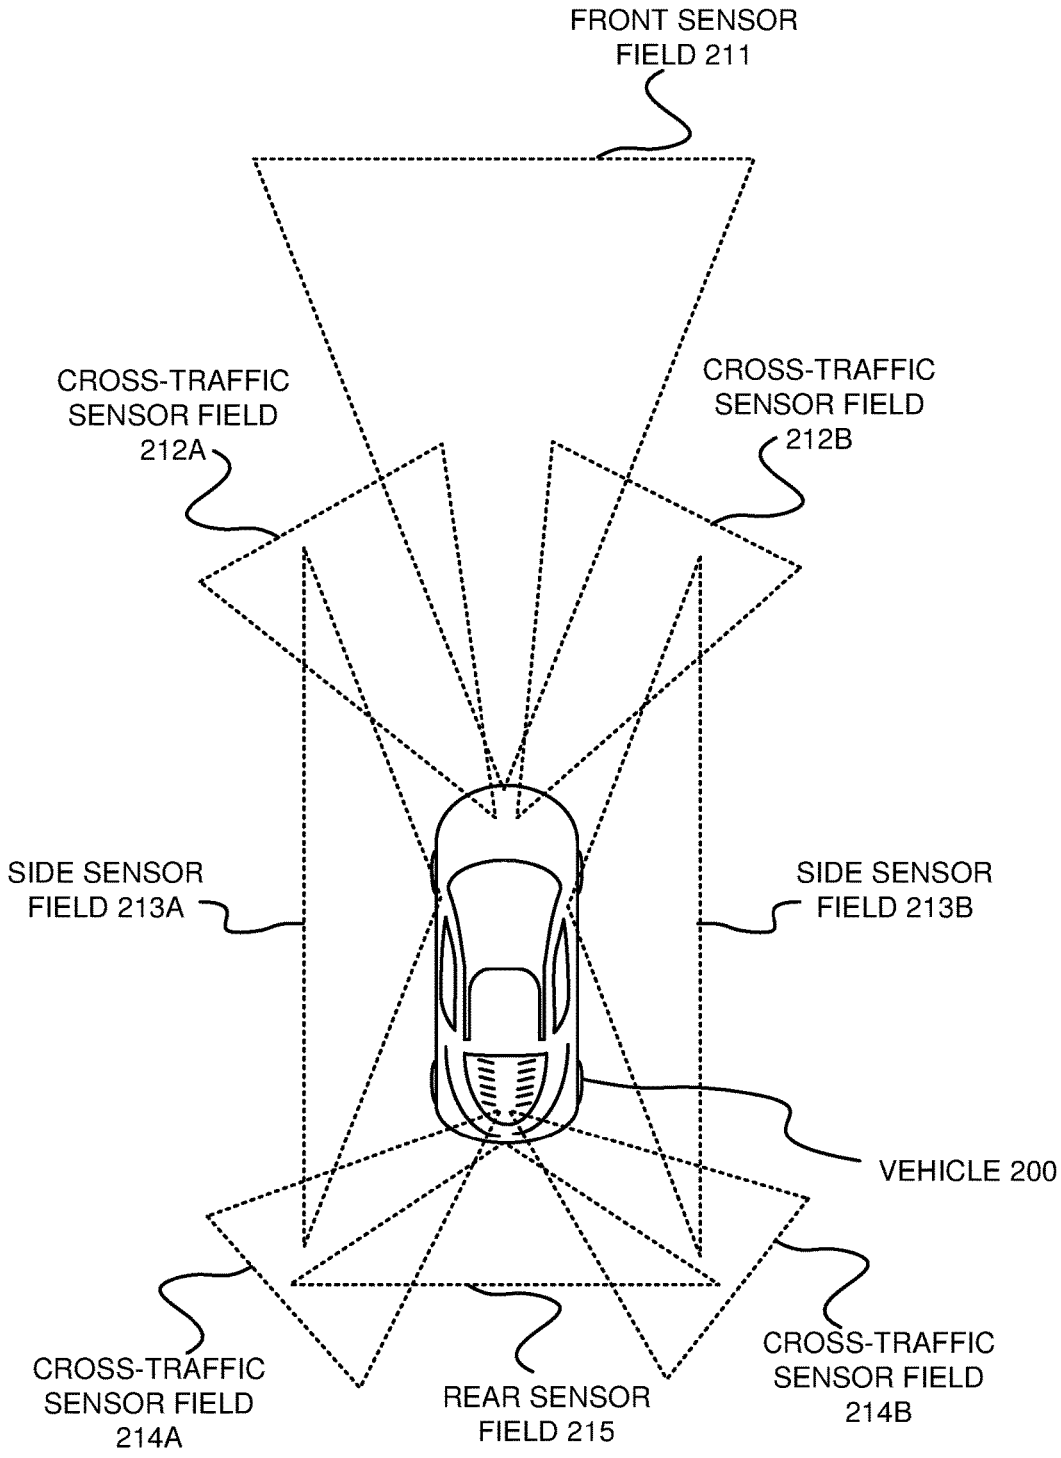

FRONT SENSOR FIELD 211
CROSS-TRAFFIC SENSOR FIELD 212A
CROSS-TRAFFIC SENSOR FIELD 212B
SIDE SENSOR FIELD 213A
SIDE SENSOR FIELD 213B
VEHICLE 200
CROSS-TRAFFIC SENSOR FIELD 214A
REAR SENSOR FIELD 215
CROSS-TRAFFIC SENSOR FIELD 214B

FIG. 2A

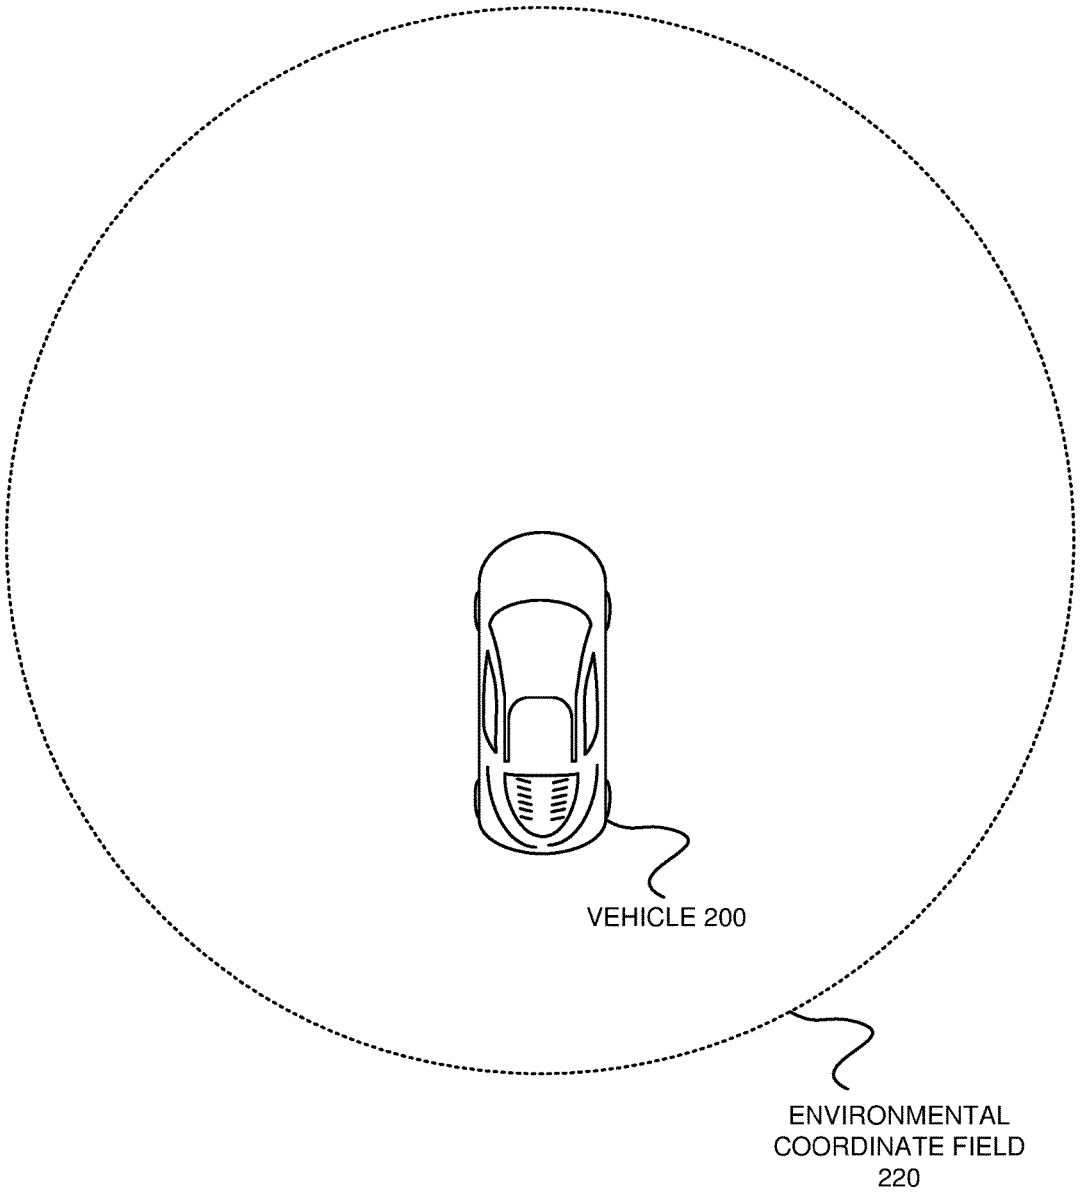


FIG. 2B

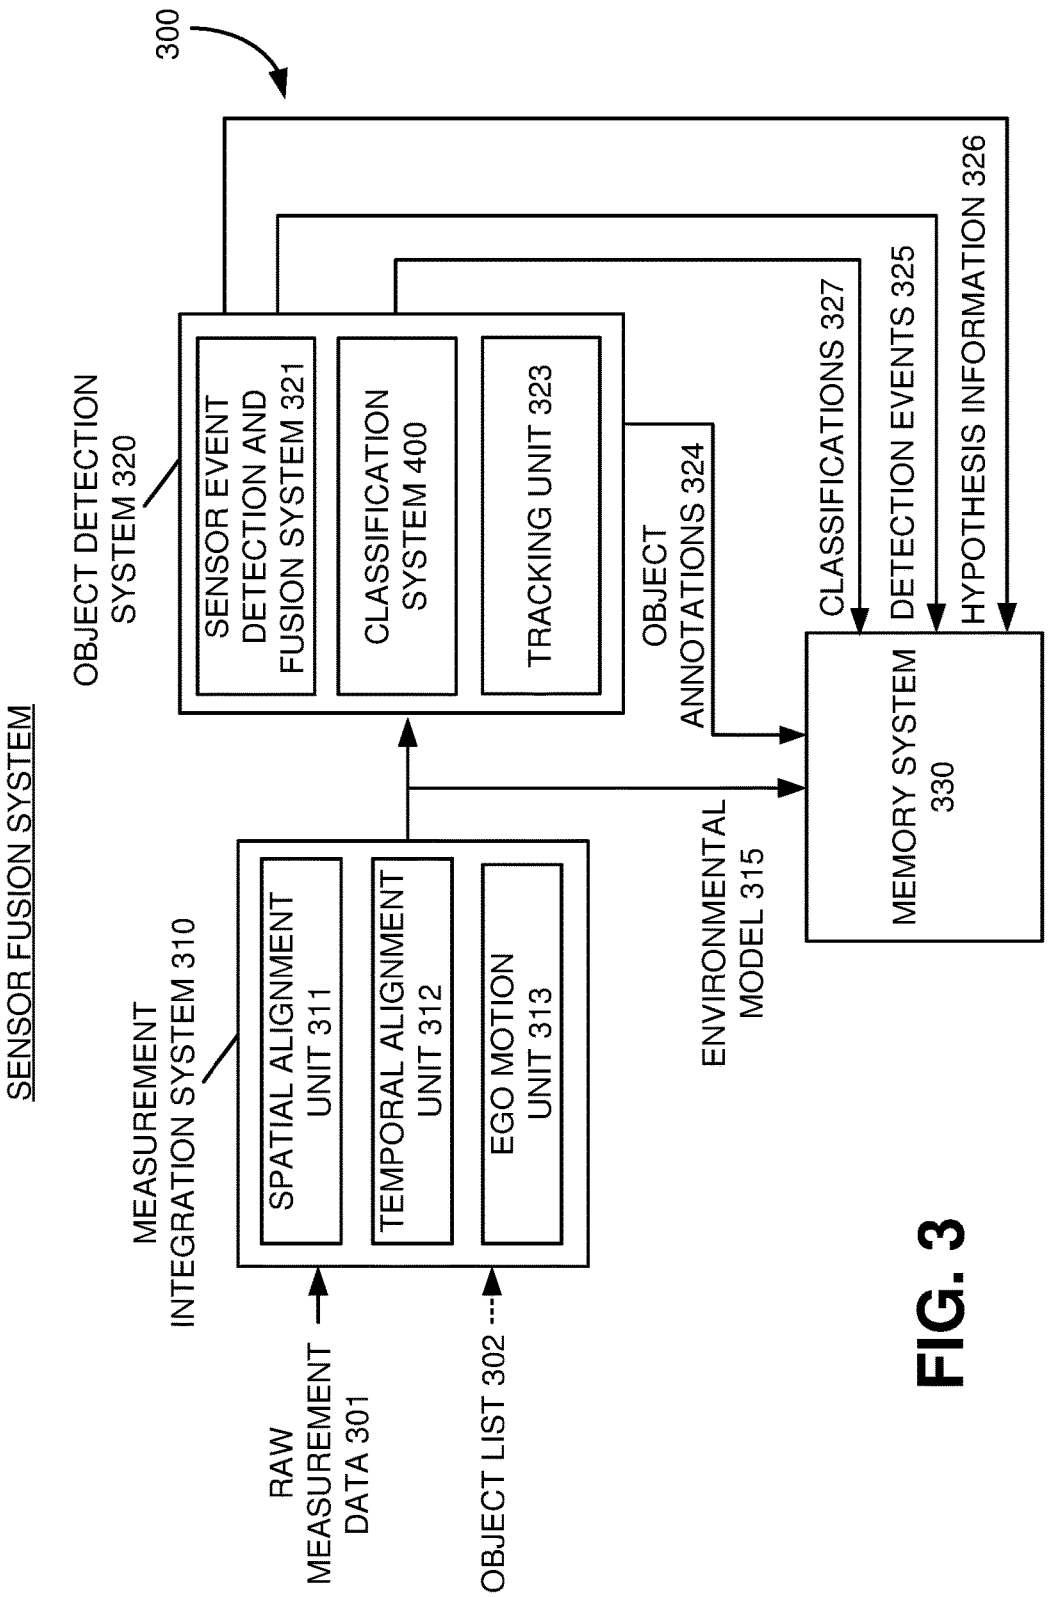

300
SENSOR FUSION SYSTEM
OBJECT DETECTION SYSTEM 320
SENSOR EVENT DETECTION AND FUSION SYSTEM 321
CLASSIFICATION SYSTEM 400
TRACKING UNIT 323
OBJECT ANNOTATIONS 324
CLASSIFICATIONS 327
DETECTION EVENTS 325
HYPOTHESIS INFORMATION 326
MEMORY SYSTEM 330
ENVIRONMENTAL MODEL 315
MEASUREMENT INTEGRATION SYSTEM 310
SPATIAL ALIGNMENT UNIT 311
TEMPORAL ALIGNMENT UNIT 312
EGO MOTION UNIT 313
RAW MEASUREMENT DATA 301
OBJECT LIST 302

FIG. 3

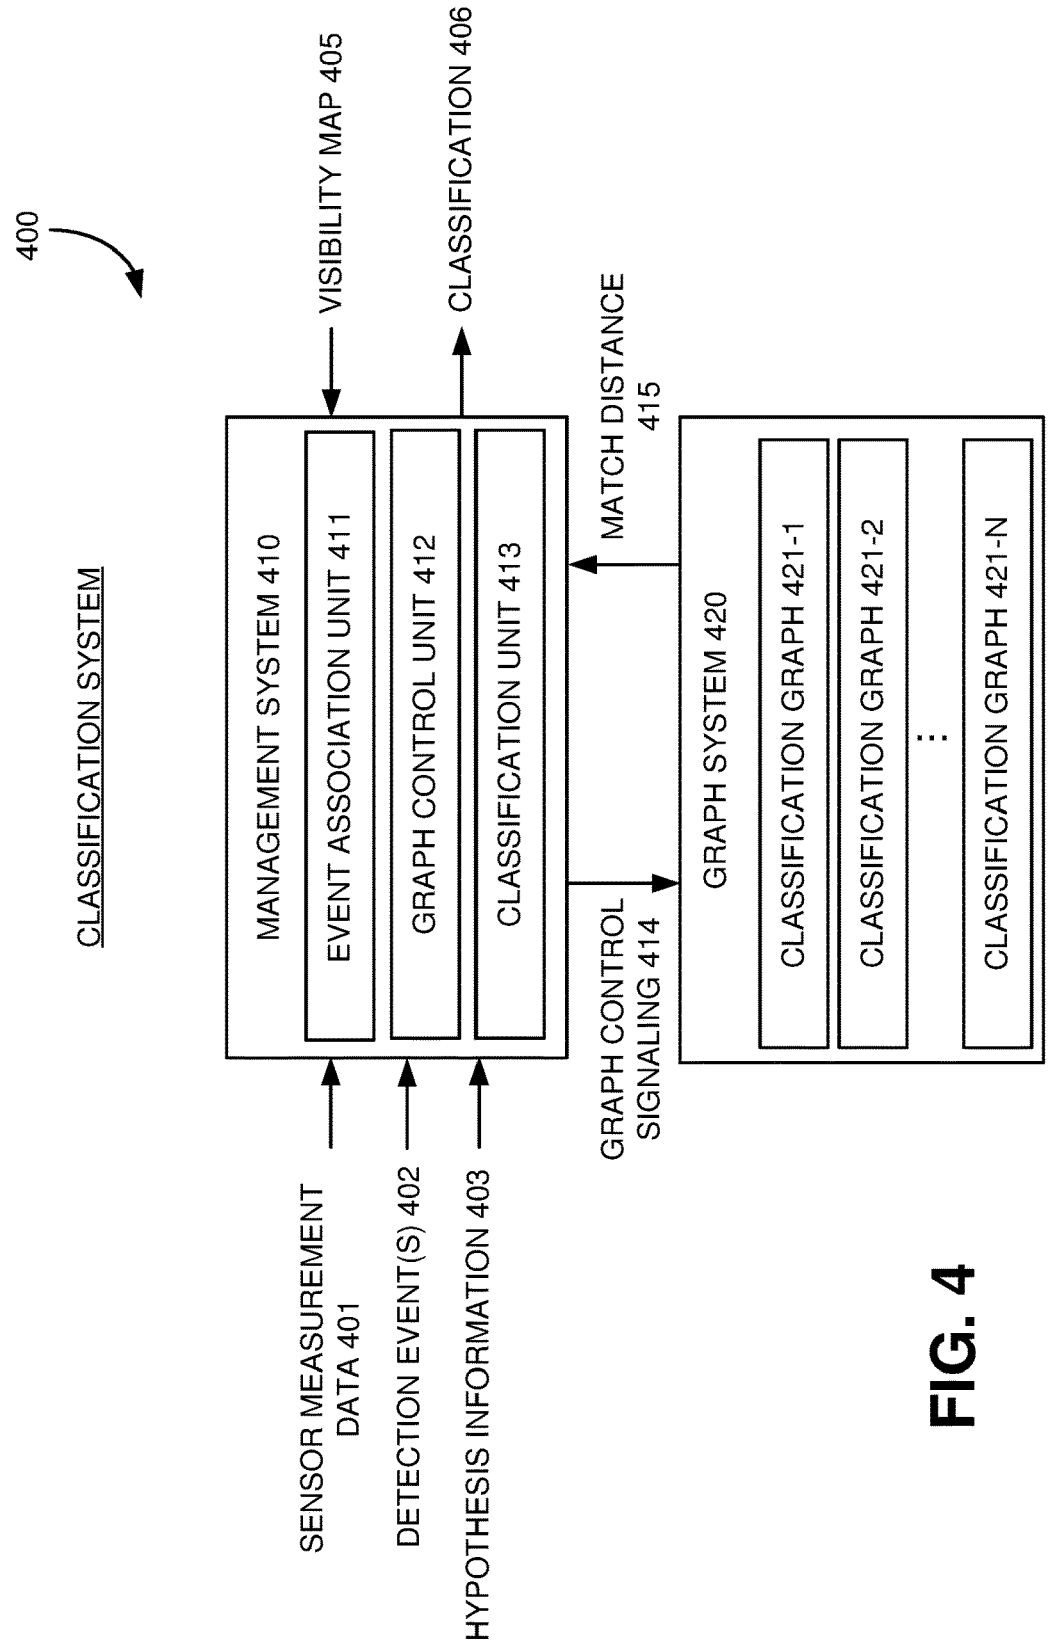

400
CLASSIFICATION SYSTEM
VISIBILITY MAP 405
CLASSIFICATION 406
MANAGEMENT SYSTEM 410
EVENT ASSOCIATION UNIT 411
GRAPH CONTROL UNIT 412
CLASSIFICATION UNIT 413
MATCH DISTANCE 415
GRAPH SYSTEM 420
CLASSIFICATION GRAPH 421-1
CLASSIFICATION GRAPH 421-2
CLASSIFICATION GRAPH 421-N
GRAPH CONTROL SIGNALING 414
SENSOR MEASUREMENT DATA 401
DETECTION EVENT(S) 402
HYPOTHESIS INFORMATION 403

FIG. 4

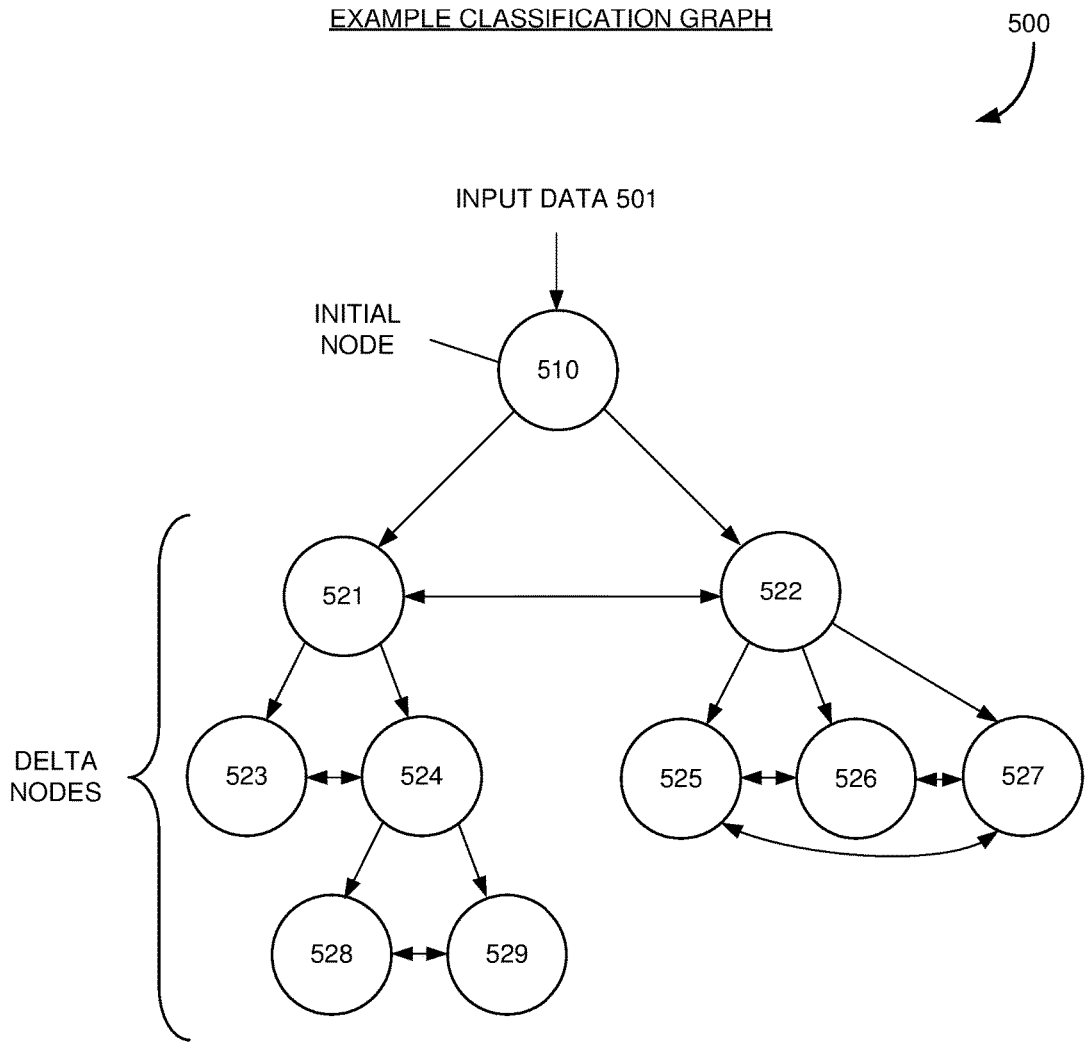

EXAMPLE CLASSIFICATION GRAPH
500
INPUT DATA 501
INITIAL NODE
510
521
522
523
524
525
526
527
528
529
DELTA NODES

FIG. 5A

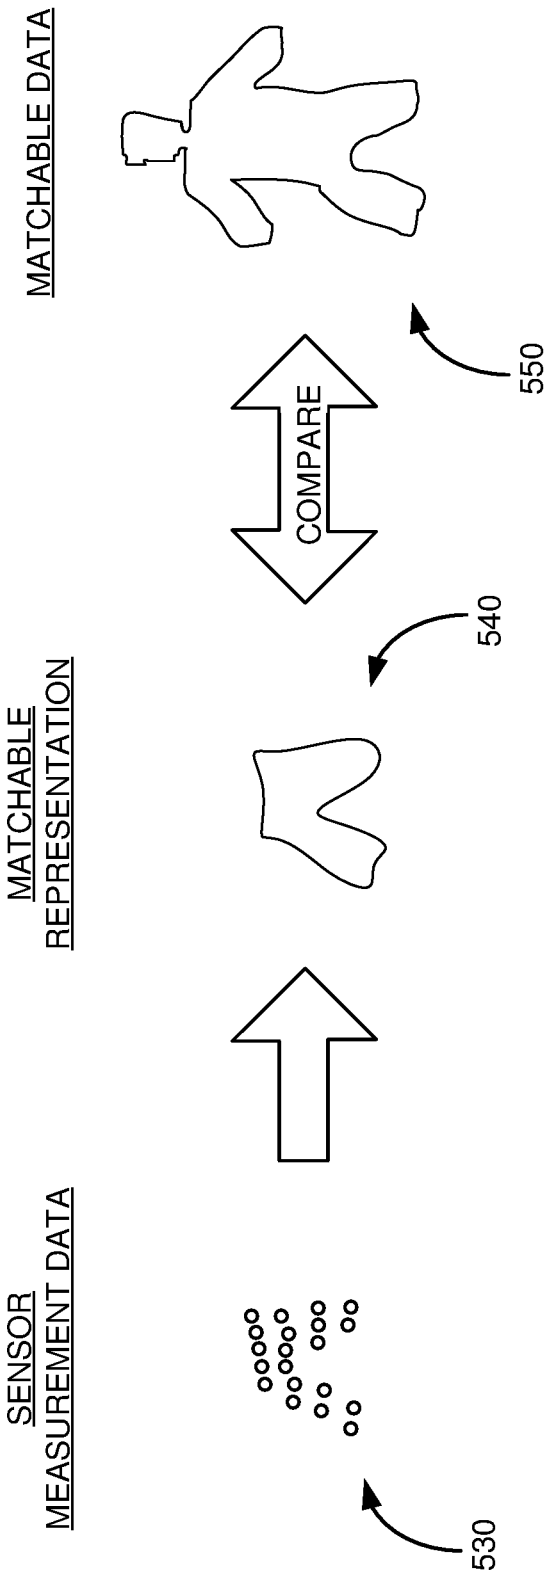

MATCHABLE DATA
550
COMPARE
MATCHABLE REPRESENTATION
540
SENSOR MEASUREMENT DATA
530

FIG. 5B

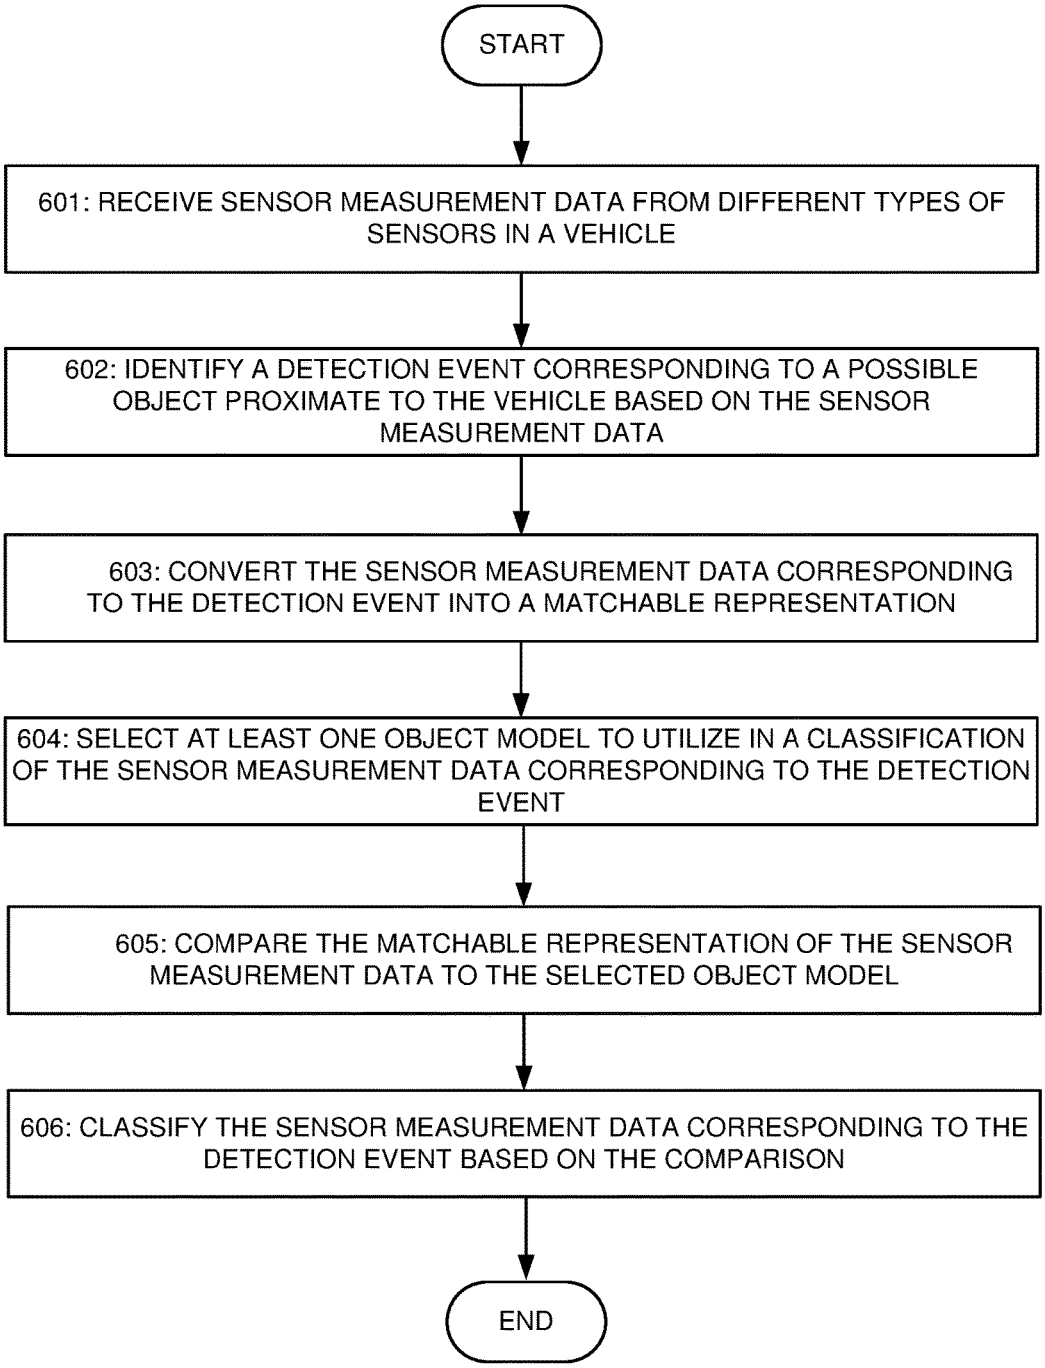

START
601: RECEIVE SENSOR MEASUREMENT DATA FROM DIFFERENT TYPES OF SENSORS IN A VEHICLE
602: IDENTIFY A DETECTION EVENT CORRESPONDING TO A POSSIBLE OBJECT PROXIMATE TO THE VEHICLE BASED ON THE SENSOR MEASUREMENT DATA
603: CONVERT THE SENSOR MEASUREMENT DATA CORRESPONDING TO THE DETECTION EVENT INTO A MATCHABLE REPRESENTATION
604: SELECT AT LEAST ONE OBJECT MODEL TO UTILIZE IN A CLASSIFICATION OF THE SENSOR MEASUREMENT DATA CORRESPONDING TO THE DETECTION EVENT
605: COMPARE THE MATCHABLE REPRESENTATION OF THE SENSOR MEASUREMENT DATA TO THE SELECTED OBJECT MODEL
606: CLASSIFY THE SENSOR MEASUREMENT DATA CORRESPONDING TO THE DETECTION EVENT BASED ON THE COMPARISON
END

FIG. 6

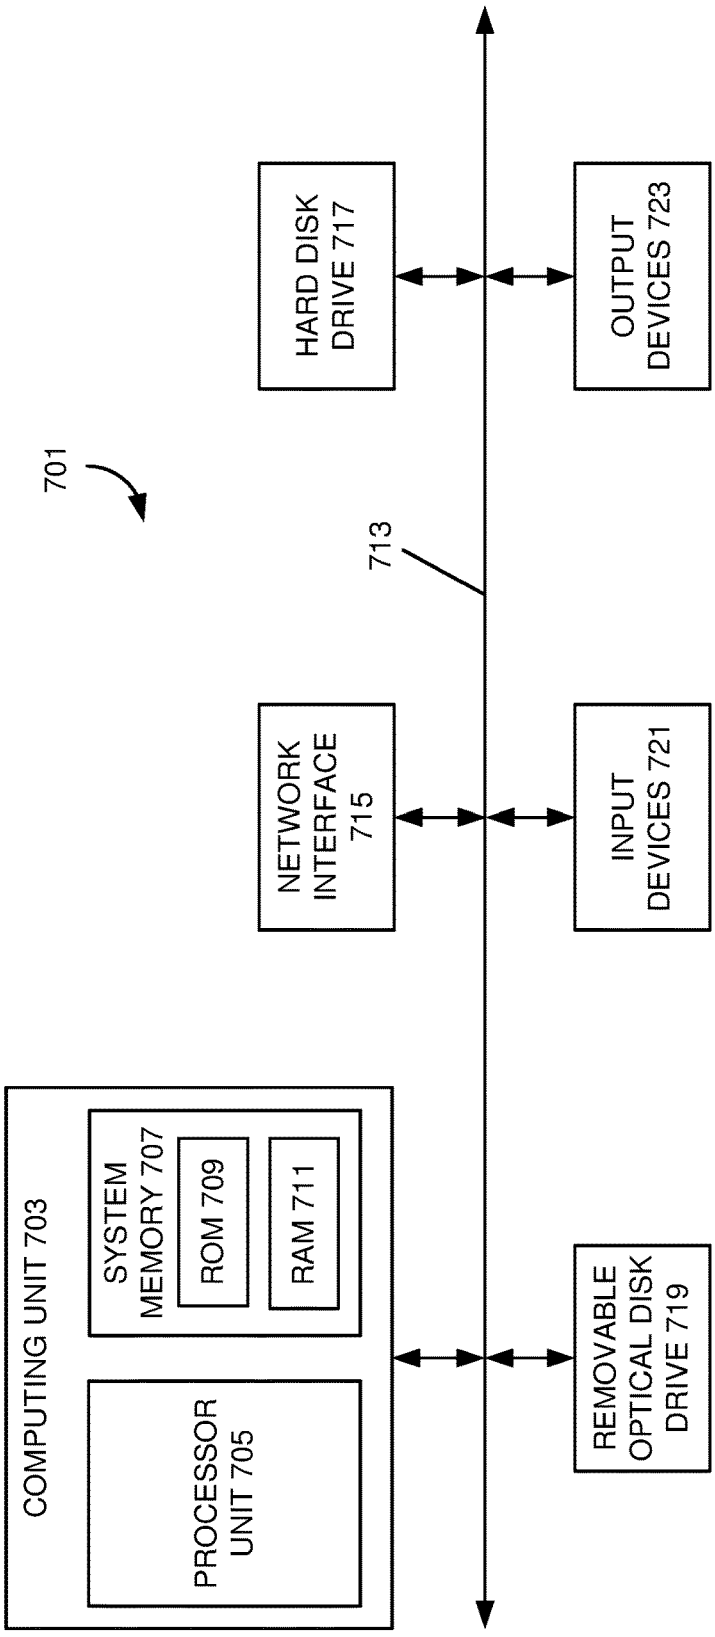

701
713
HARD DISK DRIVE 717
OUTPUT DEVICES 723
NETWORK INTERFACE 715
INPUT DEVICES 721
REMOVABLE OPTICAL DISK DRIVE 719
COMPUTING UNIT 703
PROCESSOR UNIT 705
SYSTEM MEMORY 707
ROM 709
RAM 711

FIG. 7

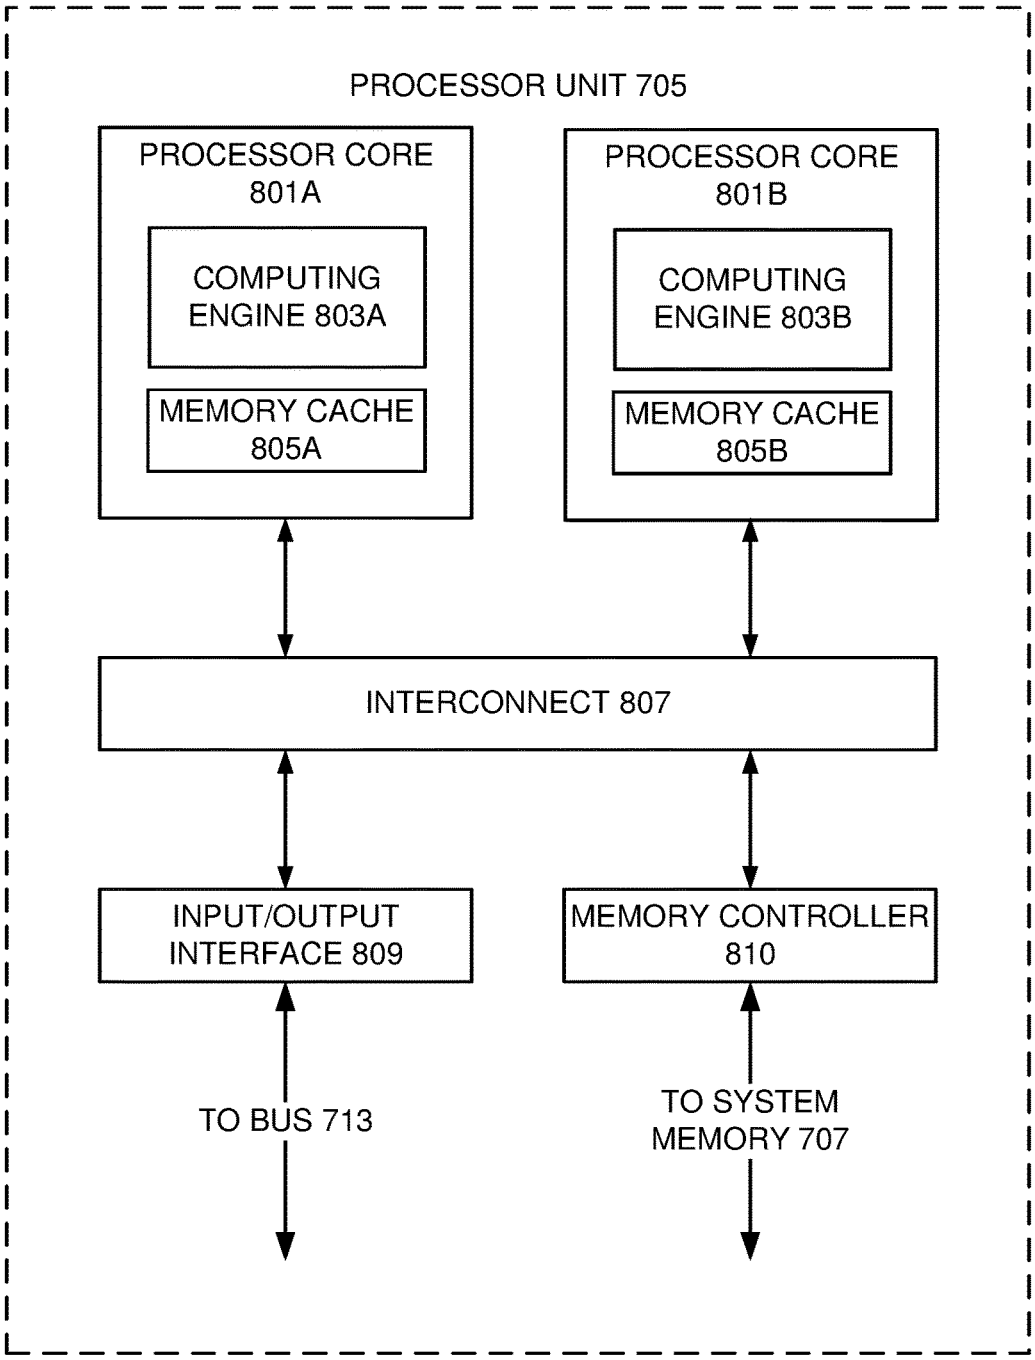

PROCESSOR UNIT 705
PROCESSOR CORE 801A
COMPUTING ENGINE 803A
MEMORY CACHE 805A
PROCESSOR CORE 801B
COMPUTING ENGINE 803B
MEMORY CACHE 805B
INTERCONNECT 807
INPUT/OUTPUT INTERFACE 809
MEMORY CONTROLLER 810
TO BUS 713
TO SYSTEM MEMORY 707

FIG. 8

EMBEDDED AUTOMOTIVE PERCEPTION WITH MACHINE LEARNING CLASSIFICATION OF SENSOR DATA

RELATED APPLICATION

[0001] This patent application claims priority to U.S. Provisional Patent Application No. 62/492,416, filed May 1, 2017, which is incorporated by reference herein.

TECHNICAL FIELD

[0002] This application is generally related to automated driving and assistance systems and, more specifically, to embedded automotive perception with machine learning classification of sensor data.

BACKGROUND

[0003] Many modern vehicles include built-in advanced driver assistance systems (ADAS) to provide automated safety and/or assisted driving functionality. For example, these advanced driver assistance systems can have applications to implement adaptive cruise control, automatic parking, automated braking, blind spot monitoring, collision avoidance, driver drowsiness detection, lane departure warning, or the like. The next generation of vehicles can include autonomous driving (AD) systems to control and navigate the vehicles independent of human interaction.

[0004] These vehicles typically include multiple sensors, such as one or more cameras, a Light Detection and Ranging (LIDAR) sensor, a Radio Detection and Ranging (RADAR) system, ultrasonic, or the like, to measure different portions of the environment around the vehicles. Each sensor processes their own measurements captured over time to detect an object within their field of view, and then provide a list of detected objects to an application in the advanced driver assistance systems or the autonomous driving systems to which the sensor is dedicated. In some instances, the sensors can also provide a confidence level corresponding to their detection of objects on the list based on their captured measurements.

[0005] The applications in the advanced driver assistance systems or the autonomous driving systems can utilize the list of objects received from their corresponding sensors and, in some cases, the associated confidence levels of their detection, to implement automated safety and/or driving functionality. For example, when a RADAR sensor in the front of a vehicle provides the advanced driver assistance system in the vehicle a list having an object in a current path of the vehicle, the application corresponding to front-end collision in the advanced driver assistance system can provide a warning to the driver of the vehicle or control vehicle in order to avoid a collision with the object.

[0006] Because each application has dedicated sensors, the application can receive a list of objects from the dedicated sensors that provides the application a fixed field of view in around a portion of the vehicle. When multiple sensors for an application have at least partially overlapping fields of view, the application can integrate object lists from its multiple dedicated sensors for the fixed field of view around the portion of the vehicle for the application. Since the vehicle moves, however, having a narrow field of view provided from the sensors can leave the application blind to potential objects. Conversely, widening the field of view can increase cost, for example, due to additional sensors, and add data processing latency.

SUMMARY

[0007] This application discloses a computing system to perform machine learning classification of sensor measurement data in an assisted or automated driving system of a vehicle. The computing system can receive sensor measurement data from multiple different sensor modalities, which the computing system can temporally and spatially align into an environmental model. The computing system can detect events, such as data point clusters or image features, in the sensor measurement data.

[0008] The computing system can implement a machine learning object classifier to generate a classification for the sensor measurement data corresponding to the detection events, and optionally a confidence level associated with the generated classification. The computing system can input a matchable representation of the sensor measurement data into the machine learning object classifier, which can compare the matchable representation of the sensor measurement data to at least one object model describing a type of object capable of being located proximate to the vehicle. The object model can include matchable data for a certain object type, which can have various poses, orientations, transitional states, potential deformations for the poses or orientations, textural features, or the like, to be compared against the matchable representation. Based on the comparison, the computing system can classify (or pre-classify) the sensor measurement data as corresponding to an object type.

[0009] In some embodiments, the computing system can generate the matchable representation of the sensor measurement data by determining a boundary outline for the sensor measurement data and creating surface area data within the boundary outline based, at least in part, on the sensor measurement data. The computing system also can estimate distances between the vehicle and objects associated with the sensor measurement data in the detection events, which can be utilized by the machine learning object classifier in the comparison between the matchable representation of the sensor measurement data and the object model.

[0010] In some embodiments, the machine learning object classifier can include multiple classification graphs or computational graphs, for example, each to describe a different object model. In some embodiments, a classification graph can include multiple nodes, each configured to include matchable data corresponding to a subset of the various poses, orientations, transitional states, potential deformations, textural features, in the object model. The computing system implementing the machine learning object classifier can perform the comparison of the matchable representation of the sensor measurement data to the object model by traversing the nodes in one or more of the classification graphs to classify to the sensor measurement data. In some embodiments, the classification system also can include a management node to select which of the classification graphs to traverse during the comparison and to control the traversal through the selected classification graphs, for example, based on outputs from the nodes of the selected classification graphs.

[0011] The computing system can utilize the classification of the sensor measurement data and associated confidence

levels in a variety of ways. In some embodiments, the computing system can utilize the classification as a pre-classification for subsequent tracking of objects in the environment model. A control system for the vehicle can be configured to control operation of the vehicle based, at least in part, on the classified type of the object for the sensor measurement data. In some embodiments, the classification can identify specific information around the vehicle, such as whether an adjacent vehicle has brake lights illuminated, a color of a stop light, or the like, which can be utilized by situational awareness functionality of the computing system. Embodiments will be described below in greater detail.

DESCRIPTION OF THE DRAWINGS

[0012] FIG. **1** illustrates an example autonomous driving system according to various embodiments.

[0013] FIG. **2**A illustrates an example measurement coordinate fields for a sensor system deployed in a vehicle according to various embodiments.

[0014] FIG. **2**B illustrates an example environmental coordinate field associated with an environmental model for a vehicle according to various embodiments.

[0015] FIG. **3** illustrates an example sensor fusion system according to various examples.

[0016] FIG. **4** illustrates an example classification system in a sensor fusion system according to various embodiments.

[0017] FIG. **5**A illustrates an example classification graph in a machine learning object classifier implementation of a classification system according to various examples.

[0018] FIG. **5**B illustrates an example flow for comparing sensor measurement data to matchable data in a node of classification graph according to various embodiments.

[0019] FIG. **6** illustrates an example flowchart for classification of sensor measurement data according to various examples.

[0020] FIGS. **7** and **8** illustrate an example of a computer system of the type that may be used to implement various embodiments of the invention.

DETAILED DESCRIPTION

Sensor Fusion for Autonomous Driving

[0021] FIG. **1** illustrates an example autonomous driving system **100** according to various embodiments. Referring to FIG. **1**, the autonomous driving system **100**, when installed in a vehicle, can sense an environment surrounding the vehicle and control operation of the vehicle based, at least in part, on the sensed environment.

[0022] The autonomous driving system **100** can include a sensor system **110** having multiple sensors, each of which can measure different portions of the environment surrounding the vehicle and output the measurements as raw measurement data **115**. The raw measurement data **115** can include characteristics of light, electromagnetic waves, or sound captured by the sensors, such as an intensity or a frequency of the light, electromagnetic waves, or the sound, an angle of reception by the sensors, a time delay between a transmission and the corresponding reception of the light, electromagnetic waves, or the sound, a time of capture of the light, electromagnetic waves, or sound, or the like.

[0023] The sensor system **110** can include multiple different types of sensors, such as an image capture device **111**, a Radio Detection and Ranging (RADAR) device **112**, a Light Detection and Ranging (LIDAR) device **113**, an ultra-sonic device **114**, one or more microphones, infrared or night-vision cameras, time-of-flight cameras, cameras capable of detecting and transmitting differences in pixel intensity, or the like. The image capture device **111**, such as one or more cameras or event-based cameras, can capture at least one image of at least a portion of the environment surrounding the vehicle. The image capture device **111** can output the captured image(s) as raw measurement data **115**, which, in some embodiments, can be unprocessed and/or uncompressed pixel data corresponding to the captured image(s).

[0024] The RADAR device **112** can emit radio signals into the environment surrounding the vehicle. Since the emitted radio signals may reflect off of objects in the environment, the RADAR device **112** can detect the reflected radio signals incoming from the environment. The RADAR device **112** can measure the incoming radio signals by, for example, measuring a signal strength of the radio signals, a reception angle, a frequency, or the like. The RADAR device **112** also can measure a time delay between an emission of a radio signal and a measurement of the incoming radio signals from the environment that corresponds to emitted radio signals reflected off of objects in the environment. The RADAR device **112** can output the measurements of the incoming radio signals as the raw measurement data **115**.

[0025] The LIDAR device **113** can transmit light, such as from a laser or other optical transmission device, into the environment surrounding the vehicle. The transmitted light, in some embodiments, can be pulses of ultraviolet light, visible light, near infrared light, or the like. Since the transmitted light can reflect off of objects in the environment, the LIDAR device **113** can include a photo detector to measure light incoming from the environment. The LIDAR device **113** can measure the incoming light by, for example, measuring an intensity of the light, a wavelength, or the like. The LIDAR device **113** also can measure a time delay between a transmission of a light pulse and a measurement of the light incoming from the environment that corresponds to the transmitted light having reflected off of objects in the environment. The LIDAR device **113** can output the measurements of the incoming light and the time delay as the raw measurement data **115**.

[0026] The ultra-sonic device **114** can emit acoustic pulses, for example, generated by transducers or the like, into the environment surrounding the vehicle. The ultra-sonic device **114** can detect ultra-sonic sound incoming from the environment, such as, for example, the emitted acoustic pulses having been reflected off of objects in the environment. The ultra-sonic device **114** also can measure a time delay between emission of the acoustic pulses and reception of the ultra-sonic sound from the environment that corresponds to the emitted acoustic pulses having reflected off of objects in the environment. The ultra-sonic device **114** can output the measurements of the incoming ultra-sonic sound and the time delay as the raw measurement data **115**.

[0027] The different sensors in the sensor system **110** can be mounted in the vehicle to capture measurements for different portions of the environment surrounding the vehicle. FIG. **2**A illustrates an example measurement coordinate fields for a sensor system deployed in a vehicle **200** according to various embodiments. Referring to FIG. **2**A, the vehicle **200** can include multiple different sensors capable of detecting incoming signals, such as light signals, electromagnetic signals, and sound signals. Each of these

different sensors can have a different field of view into an environment around the vehicle **200**. These fields of view can allow the sensors to measure light and/or sound in different measurement coordinate fields.

[0028] The vehicle in this example includes several different measurement coordinate fields, including a front sensor field **211**, multiple cross-traffic sensor fields **212**A, **212**B, **214**A, and **214**B, a pair of side sensor fields **213**A and **213**B, and a rear sensor field **215**. Each of the measurement coordinate fields can be sensor-centric, meaning that the measurement coordinate fields can describe a coordinate region relative to a location of its corresponding sensor.

[0029] Referring back to FIG. **1**, the autonomous driving system **100** can include a sensor fusion system **300** to receive the raw measurement data **115** from the sensor system **110** and populate an environmental model **121** associated with the vehicle with the raw measurement data **115**. In some embodiments, the environmental model **121** can have an environmental coordinate field corresponding to a physical envelope surrounding the vehicle, and the sensor fusion system **300** can populate the environmental model **121** with the raw measurement data **115** based on the environmental coordinate field. In some embodiments, the environmental coordinate field can be a non-vehicle centric coordinate field, for example, a world coordinate system, a path-centric coordinate field, a coordinate field parallel to a road surface utilized by the vehicle, or the like.

[0030] FIG. **2**B illustrates an example environmental coordinate field **220** associated with an environmental model for the vehicle **200** according to various embodiments. Referring to FIG. **2**B, an environment surrounding the vehicle **200** can correspond to the environmental coordinate field **220** for the environmental model. The environmental coordinate field **220** can be vehicle-centric and provide a 360 degree area around the vehicle **200**. The environmental model can be populated and annotated with information detected by the sensor fusion system **300** or inputted from external sources. Embodiments will be described below in greater detail.

[0031] Referring back to FIG. **1**, to populate the raw measurement data **115** into the environmental model **121** associated with the vehicle, the sensor fusion system **300** can spatially align the raw measurement data **115** to the environmental coordinate field of the environmental model **121**. The sensor fusion system **300** also can identify when the sensors captured the raw measurement data **115**, for example, by time stamping the raw measurement data **115** when received from the sensor system **110**. The sensor fusion system **300** can populate the environmental model **121** with the time stamp or other time-of-capture information, which can be utilized to temporally align the raw measurement data **115** in the environmental model **121**. In some embodiments, the sensor fusion system **300** can analyze the raw measurement data **115** from the multiple sensors as populated in the environmental model **121** to detect a sensor event or at least one object in the environmental coordinate field associated with the vehicle. The sensor event can include a sensor measurement event corresponding to a presence of the raw measurement data **115** in the environmental model **121**, for example, above a noise threshold. The sensor event can include a sensor detection event corresponding to a spatial and/or temporal grouping of the raw measurement data **115** in the environmental model **121**. The object can correspond to spatial grouping of the raw measurement data **115** having been tracked in the environmental model **121** over a period of time, allowing the sensor fusion system **300** to determine the raw measurement data **115** corresponds to an object around the vehicle. The sensor fusion system **300** can populate the environment model **121** with an indication of the detected sensor event or detected object and a confidence level of the detection. Embodiments of sensor fusion and sensor event detection or object detection will be described below in greater detail.

[0032] The sensor fusion system **300**, in some embodiments, can generate feedback signals **116** to provide to the sensor system **110**. The feedback signals **116** can be configured to prompt the sensor system **110** to calibrate one or more of its sensors. For example, the sensor system **110**, in response to the feedback signals **116**, can re-position at least one of its sensors, expand a field of view of at least one of its sensors, change a refresh rate or exposure time of at least one of its sensors, alter a mode of operation of at least one of its sensors, or the like.

[0033] The autonomous driving system **100** can include a driving functionality system **120** to receive at least a portion of the environmental model **121** from the sensor fusion system **300**. The driving functionality system **120** can analyze the data included in the environmental model **121** to implement automated driving functionality or automated safety and assisted driving functionality for the vehicle. The driving functionality system **120** can generate control signals **131** based on the analysis of the environmental model **121**.

[0034] The autonomous driving system **100** can include a vehicle control system **130** to receive the control signals **131** from the driving functionality system **120**. The vehicle control system **130** can include mechanisms to control operation of the vehicle, for example by controlling different functions of the vehicle, such as braking, acceleration, steering, parking brake, transmission, user interfaces, warning systems, or the like, in response to the control signals.

[0035] FIG. **3** illustrates an example sensor fusion system **300** according to various examples. Referring to FIG. **3**, the sensor fusion system **300** can include a measurement integration system **310** to receive raw measurement data **301** from multiple sensors mounted in a vehicle. The measurement integration system **310** can generate an environmental model **315** for the vehicle, which can be populated with the raw measurement data **301**.

[0036] The measurement integration system **310** can include a spatial alignment unit **311** to correlate measurement coordinate fields of the sensors to an environmental coordinate field for the environmental model **315**. The measurement integration system **310** can utilize this correlation to convert or translate locations for the raw measurement data **301** within the measurement coordinate fields into locations within the environmental coordinate field. The measurement integration system **310** can populate the environmental model **315** with the raw measurement data **301** based on the correlation between the measurement coordinate fields of the sensors to the environmental coordinate field for the environmental model **315**.

[0037] The measurement integration system **310** also can temporally align the raw measurement data **301** from different sensors in the sensor system. In some embodiments, the measurement integration system **310** can include a temporal alignment unit **312** to assign time stamps to the raw measurement data **301** based on when the sensor captured the raw measurement data **301**, when the raw measurement

data **301** was received by the measurement integration system **310**, or the like. In some embodiments, the temporal alignment unit **312** can convert a capture time of the raw measurement data **301** provided by the sensors into a time corresponding to the sensor fusion system **300**. The measurement integration system **310** can annotate the raw measurement data **301** populated in the environmental model **315** with the time stamps for the raw measurement data **301**. The time stamps for the raw measurement data **301** can be utilized by the sensor fusion system **300** to group the raw measurement data **301** in the environmental model **315** into different time periods or time slices. In some embodiments, a size or duration of the time periods or time slices can be based, at least in part, on a refresh rate of one or more sensors in the sensor system. For example, the sensor fusion system **300** can set a time slice to correspond to the sensor with a fastest rate of providing new raw measurement data **301** to the sensor fusion system **300**.

[0038] The measurement integration system **310** can include an ego motion unit **313** to compensate for movement of at least one sensor capturing the raw measurement data **301**, for example, due to the vehicle driving or moving in the environment. The ego motion unit **313** can estimate motion of the sensor capturing the raw measurement data **301**, for example, by utilizing tracking functionality to analyze vehicle motion information, such as global positioning system (GPS) data, inertial measurements, vehicle odometer data, video images, or the like. The tracking functionality can implement a Kalman filter, a Particle filter, optical flow-based estimator, or the like, to track motion of the vehicle and its corresponding sensors relative to the environment surrounding the vehicle.

[0039] The ego motion unit **313** can utilize the estimated motion of the sensor to modify the correlation between the measurement coordinate field of the sensor to the environmental coordinate field for the environmental model **315**. This compensation of the correlation can allow the measurement integration system **310** to populate the environmental model **315** with the raw measurement data **301** at locations of the environmental coordinate field where the raw measurement data **301** was captured as opposed to the current location of the sensor at the end of its measurement capture.

[0040] In some embodiments, the measurement integration system **310** may receive objects or object lists **302** from a variety of sources. The measurement integration system **310** can receive the object list **302** from sources external to the vehicle, such as in a vehicle-to-vehicle (V2V) communication, a vehicle-to-infrastructure (V2I) communication, a vehicle-to-pedestrian (V2P) communication, a vehicle-to-device (V2D) communication, a vehicle-to-grid (V2G) communication, or generally a vehicle-to-everything (V2X) communication. The measurement integration system **310** also can receive the objects or an object list **302** from other systems internal to the vehicle, such as from a human machine interface, mapping systems, localization system, driving functionality system, vehicle control system, or the vehicle may be equipped with at least one sensor that outputs the object list **302** rather than the raw measurement data **301**.

[0041] The measurement integration system **310** can receive the object list **302** and populate one or more objects from the object list **302** into the environmental model **315** along with the raw measurement data **301**. The object list **302** may include one or more objects, a time stamp for each object, and optionally include a spatial metadata associated with a location of objects in the object list **302**. For example, the object list **302** can include speed measurements for the vehicle, which may not include a spatial component to be stored in the object list **302** as the spatial metadata. When the object list **302** includes a confidence level associated with an object in the object list **302**, the measurement integration system **310** also can annotate the environmental model **315** with the confidence level for the object from the object list **302**.

[0042] The sensor fusion system **300** can include an object detection system **320** to receive the environmental model **315** from the measurement integration system **310**. In some embodiments, the sensor fusion system **300** can include a memory system **330** to store the environmental model **315** from the measurement integration system **310**. The object detection system **320** may access the environmental model **315** from the memory system **330**.

[0043] The object detection system **320** can analyze data stored in the environmental model **315** to detect at least one object. The sensor fusion system **300** can populate the environment model **315** with an indication of the detected object at a location in the environmental coordinate field corresponding to the detection. The object detection system **320** can identify confidence levels corresponding to the detected object, which can be based on at least one of a quantity, a quality, or a sensor diversity of raw measurement data **301** utilized in detecting the object. The sensor fusion system **300** can populate or store the confidence levels corresponding to the detected objects with the environmental model **315**. For example, the object detection system **320** can annotate the environmental model **315** with object annotations **324** or the object detection system **320** can output the object annotations **324** to the memory system **330**, which populates the environmental model **315** with the detected object and corresponding confidence level of the detection in the object annotations **324**.

[0044] The object detection system **320** can include a sensor event detection and fusion system **321** to identify detection events **325** from the data stored in the environmental model **315**. In some embodiments, the sensor event detection and fusion system **400** can identify the detection events **325** by analyzing the data stored in the environmental model **315** on a per-sensor type basis to identify patterns in the data, such as image features or data point clusters. When the sensor event detection and fusion system **321** utilizes patterns from a single sensor modality or type to generate the detection events **325**, the detection event **325** may be called a sensor detection event. In some embodiments, the sensor event detection and fusion system **321** also can associate or correlate identified patterns across multiple different sensor modalities or types to generate the detection event **325**, which can be called a fused sensor detection event.

[0045] The sensor event detection and fusion system **321** also can determine differences from adjacent frames or scans of the sensor measurement data on a per-sensor-type basis. For example, the sensor event detection and fusion system **321** can compare the received sensor measurement data from a type of sensor against sensor measurement data from a previously received frame or scan from that type of sensor to determine the differences from adjacent frames or scans of the sensor measurement data. The sensor event detection and fusion system **321** can perform this inter-frame and intra-modality comparison of the sensor measurement data

based, at least in part, on the spatial locations of the sensor measurement data in the environmental model **315**. For example, when an image capture sensor provides entire image frames, the sensor event detection and fusion system **321** can cache the entire image frames, determine inter-frame differences for the sensor measurement data from a plurality of the cached image frames. In another example, when an image capture sensor provided event-based pixels, the sensor event detection and fusion system **321** can perform pixel caching to generate an entire image from the image data. The sensor event detection and fusion system **321** can utilize the event-based pixels as the inter-frame differences in the sensor measurement data. In another example, when one or more of the RADAR sensors provides raw signal data in a frequency domain, the sensor event detection and fusion system **321** can detect one or more untracked targets from RADAR measurements. The sensor event detection and fusion system **321** can determine differences between the untracked targets in adjacent frames, which can constitute inter-frame differences in the sensor measurement data for the RADAR sensor modality.

[0046] The sensor event detection and fusion system **321** also can generate hypothesis information **326** corresponding to the detection events **325**. In some embodiments, the hypothesis information **326** can identify confidence levels corresponding to various properties or characteristics associated with the detection events **325**. The properties or characteristics associated with the detection events **325** can include unity, velocity, orientation, measurement-based virtual center of gravity, existence, size, novelty, or the like, of the sensor measurement data corresponding to the detection event **325**. The unity characteristic can identify whether the sensor measurement data corresponds to a single possible object or multiple possible objects. For example, the unity characteristic can identify when one possible object in the sensor measurement data occludes or blocks visibility to at least another possible object. The velocity characteristic can identify at least one velocity associated with the sensor measurement data. The orientation characteristic can identify a directionality of the sensor measurement data and/or an angle associated with the possible object relative to the vehicle. The measurement-based virtual center of gravity characteristic can identify a center of the possible object, for example, based on a density of the data points associated with the detection event **325**. The existence characteristic can identify whether the possible object identified by the detection event **325** is an actual object proximate to the vehicle. The size characteristic can identify or estimate a real size of the possible object associated with the detection event **325**. The novelty characteristic can identify whether the detection event **325** corresponds to a newly detected pattern or corresponding to a previous detection event **325**.

[0047] The sensor fusion system **300** can populate or store the detection events **325** and the corresponding hypothesis information **326** with the environmental model **315**. For example, the object detection system **320** can annotate the environmental model **315** with the detection events **325** and the corresponding hypothesis information **326**, or the object detection system **320** can output the detection events **325** and the corresponding hypothesis information **326** to the memory system **330**, which populates the environmental model **315** with the detection events **325** and the corresponding hypothesis information **326**.

[0048] The object detection system **320** can include a classification system **400** to classify sensor measurement data associated with the detection events **325**. In some embodiments, the classification system **400** can assign classifications **327** to the detection events **325** based on the classification of the sensor measurement data associated with the detection events **325**. The classifications **327** can correspond to a type of object associated with the detection events **325**, such as another vehicle, a pedestrian, a cyclist, an animal, a static object, or the like, which may be identified or pointed to by the hypothesis information **326**. The classifications **327** also can include a confidence level associated with the classification and/or include more specific information corresponding to a particular pose, orientation, state, or the like, of the object type. The object detection system **320** can annotate the environmental model **315** with classifications **327** or the object detection system **320** can output the classifications **327** to the memory system **330**, which populates the environmental model **315** with the classifications **327**.

[0049] The classification system **400** can implement one or more object models, each to describe a type of object capable of being located proximate to the vehicle. The object models can include matchable data for different object types, and include poses, orientations, transitional states, potential deformations for the poses or orientations, textural features, or the like, to be compared against the sensor measurement data. The classification system **400** can compare the sensor measurement data (or a modified representation of the sensor measurement data) associated with the detection events **325** to one or more of the object models, and generate the classifications **327** based on the comparison.

[0050] In some embodiments, the classification system **400** can perform the classification utilizing a machine learning object classifier. The machine learning object classifier can include multiple classification graphs or tensor graphs, for example, each to describe a different object model. In some embodiments, a classification graph can include multiple nodes, each configured to include matchable data corresponding to a subset of the various poses, orientations, transitional states, potential deformations, textural features, or the like, in the object model. The classification system **400** also can perform the classification utilizing other computational techniques, such as a feed-forward neural network, a support vector machine (SVM), or the like.

[0051] The classification system **400** can select one or more of the classification graphs based, at least in part, on the detection events **325** and/or the hypothesis information **326**. The classification system **400** can traverse nodes in the selected classification graphs, for example, with each node performing a comparison between a portion of the matchable data in the object model to the sensor measurement data (or a modified representation of the sensor measurement data). In some embodiments, the traversal of node can span across classification graphs, for example, where a node in one classification graph can directly provide information to a node in a different classification graph. Each node in the classification graphs can generate a match distance, for example, by generating a confidence level that the sensor measurement data corresponds to the matchable data of the node. The classification system **400** can utilize the match distances from the nodes to generate the classification **327** for the sensor measurement data associated with the detection events. The classification system **400** can utilize the

match distances to control the traversal through the nodes in the selected classification graphs, for example, deciding which nodes should attempt to classify the sensor measurement data and in what order, when to cease having nodes attempt to classify the sensor measurement data, or the like. Embodiments of the sensor measurement data classification will be described below in greater detail.

[0052] The object detection system **320** can include a tracking unit **323** to track the detection events **325** in the environmental model **315** over time, for example, by analyzing the annotations in the environmental model **315**, and determine whether the detection events **325** corresponds to objects in the environmental coordinate system. In some embodiments, the tracking unit **323** can utilize the classifications **327** to track the detection events **325** with at least one state change prediction model, such as a kinetic model, a probabilistic model, or other state change prediction model.

[0053] The tracking unit **323** can select the state change prediction model to utilize to track the detection events **325** based on the assigned classifications **327** of the detection events **325**. The state change prediction model may allow the tracking unit **323** to implement a state transition prediction, which can assume or predict future states of the detection events **325**, for example, based on a location of the detection events **325** in the environmental model **315**, a prior movement of the detection events **325**, a classification of the detection events **325**, or the like. In some embodiments, the tracking unit **323** implementing the kinetic model can utilize kinetic equations for velocity, acceleration, momentum, or the like, to assume or predict the future states of the detection events **325** based, at least in part, on its prior states.

[0054] The tracking unit **323** may determine a difference between the predicted future states of the detection events **325** and its actual future states, which the tracking unit **323** may utilize to determine whether the detection events **325** correspond to objects proximate to the vehicle. The tracking unit **323** can track the detection event **325** in the environmental coordinate field associated with the environmental model **315**, for example, across multiple different sensors and their corresponding measurement coordinate fields.

[0055] When the tracking unit **323**, based on the tracking of the detection events **325** with the state change prediction model, determines the detection events **325** are trackable, the tracking unit **323** can annotate the environmental model **315** to indicate the presence of trackable detect events. The tracking unit **323** can continue tracking the trackable detect events over time by implementing the state change prediction models and analyzing the environmental model **315** when updated with additional raw measurement data **301**. After annotating the environmental model **315** to indicate the presence of trackable detect events, the tracking unit **323** can continue to track the trackable detect events in the environmental coordinate field associated with the environmental model **315**, for example, across multiple different sensors and their corresponding measurement coordinate fields.

[0056] The memory system **330** can be accessible by different portions of the autonomous driving system **100**, such as a driving functionality system, which can include a localization system, a situational awareness system, or the like. The sensor fusion system **300** can produce data, such as the environmental model **315**, object annotations **324**, classifications **327**, detection events **325**, hypothesis information **326**, or the like, and portions of the driving functionality system can consume data, for example, via access to the memory system **330**. In some embodiments, the sensor fusion system **300** can include a system to output the produced data to the consuming systems based on a subscription basis or in an event-driven manner.

[0057] The consuming systems, such as the localization system or the situational awareness system, also can send requests to the sensor fusion system **300** for operations to be performed on data within the memory system **330**. For example, the sensor fusion system **300** can be requested to perform in-depth classification of a portion of the data in the memory system **330**. The sensor fusion system **300** can prompt the classification system **400** to perform the classification of the requested data and output the results to the requesting system. For example, the situational awareness system may consume data from the memory system indicating that another vehicle is in front of the vehicle including the autonomous driving system **100**, and request the sensor fusion system **300** classify the signal lights of that target vehicle, such as blinkers, brake lights, reverse lights, or the like. The sensor fusion system **300** can prompt the classification system **400** to identify locations of the signal lights in data of the target vehicle, and classify their state. Once classified by the classification system **400**, the sensor fusion system **300** can output the classification to the requesting system or add the classification to the environmental model **315**, where it is available for other vehicle systems.

Classification of Sensor Measurement Data by a Sensor Fusion System

[0058] FIG. **4** illustrates an example classification system **400** in a sensor fusion system according to various embodiments. Referring to FIG. **4**, the classification system **400** can include a management system **410** and a graph system **420**, which can operate in conjunction to generate a classification **406** for sensor measurement data **401**. The classification **406** can identify to a type of object associated with the sensor measurement data **401**, such as another vehicle, a pedestrian, a cyclist, an animal, a static object, or the like. The classification **406** also can include a confidence level associated with the identification of the object type and/or include more specific information corresponding to a particular pose, orientation, state, or the like, of the object type.

[0059] The graph system **420** can define one or more object models describing different types of objects capable of being located proximate to the vehicle, such as another vehicle, a pedestrian, a cyclist, an animal, a static object, or the like. The object models can include matchable data for different object types, and include poses, orientations, transitional states, potential deformations for the poses or orientations, textural features, or the like, to be compared against the sensor measurement data.

[0060] The management system **410** can select one or more of the object models to utilize in the generation of the classification **406** for the sensor measurement data **401**. The management system **410** can prepare the sensor measurement data **401** for comparison to the selected object models, and direct the graph system **420** via graph control signaling **414** to apply the selected object models to the prepared version of the sensor measurement data **401**. The graph system **420** can generate at least one match distance **415** based on the application of the prepared version of the sensor measurement data **401** to the selected object models,

and the management system **410** can generate the classification **406** for the sensor measurement data **401** based, at least in part, on the match distance **415** from the graph system **420**.

[0061] The management system **410** can receive the sensor measurement data **401**, detection events **402**, and hypothesis information **403**, for example, included in an environmental model for the vehicle including the sensor fusion system. The sensor measurement data **401** can be collected from multiple different sensor modalities and can temporally and spatially align into the environmental model. The detection events **402** and the hypothesis information **403** can describe properties or characteristics of at least a portion of the sensor measurement data **401**.

[0062] The management system **410**, in some embodiments, can generate a matchable representation of the sensor measurement data **401**. The matchable representation of the sensor measurement data **401** can be based on data captured from any or all sensor modalities. The matchable representation of the sensor measurement data **401** can include a two-dimensional shape description of the sensor measurement data **401**, three-dimensional structure description of the sensor measurement data **401**, a textural description of the sensor measurement data **401**, a geometric description of the sensor measurement data **401**, a structural or skeletonized description of the sensor measurement data **401**, include reflectivity features, or the like. In some embodiments, the management system **410** can perform skeleton extraction on the sensor measurement data **401**, which can divide an object in the sensor measurement data **401** into different segments and include a connectivity graph to identify how the different segments were linked in the sensor measurement data **401**. For example, the management system **410** can generate a shape outline or structural body corresponding to the sensor measurement data **401**, and thin the shape outline or structural body into lines. The management system **410** can identify points associated with the lines, such as endpoint, vertices where multiple lines connect, or the like. The management system **410** can generate the skeleton of the sensor measurement data **401** from the lines, the identified points, and relationships between the lines or connectivity of the lines. In some embodiments, the management system **410** can generate the lines directly from the sensor measurement data **401**, for example, generate lines connecting data points in the sensor measurement data **401**.

[0063] The management system **410** also can extract geometric primitives, such as U-shapes, L-shapes, curves, or the like, from the sensor measurement data **401**, and include them in the matchable representation of the sensor measurement data **401**. In some embodiments, the management system can determine a boundary outline for the sensor measurement data **401** and creating surface area data within the boundary outline based, at least in part, on the sensor measurement data **401**. For example, the management system **410** can close or fill-in the surface area of the boundary outline by expanding existing data points in the sensor measurement data **401**, connecting existing data points in the sensor measurement data **401**, or the like. The management system **410** also can identify textural features from the sensor measurement data **401**, for example, by utilizing a Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), frequency domain features, Haar-like features using Haar wavelets, or the like, to extract the textural features from the sensor measurement data **401**.

[0064] The matchable representation of the sensor measurement data **401** can be a two-dimensional representation of the sensor measurement data **401** or a three-dimensional representation of the sensor measurement data **401**. When generating the three-dimensional representation of the sensor measurement data **401**, certain volumetric aspects of the matchable representation may not be observable from the sensor measurement data **401**. In these cases, the management system **410** may estimate the three-dimensional structure associated with the sensor measurement data **401** based, at least in part, on symmetry in the sensor measurement data **401**, completeness of the sensor measurement data **401**, curvature sensor measurement data **401**, the hypothesis information **403**, or the like.

[0065] The management system **410** also can estimate distances between the vehicle and objects associated with the sensor measurement data **401**. In some embodiments, the sensor measurement data **401** can include a distance measurement in the data points, or the management system **410** can ascertain the distance based on the spatial location of the sensor measurement data **401** in the environmental model.

[0066] The management system **410** can include an event association unit **411** to determine whether one of the detection events **402** corresponds to one or more detection events **402** previously received by the management system **410**. In some embodiments, the event association unit **411** can associate the detection event **402** with the one or more previously received detection events **402** based, at least in part, on their spatial locations relative to each other, sensor update rates, types of sensor measurement data correlated to the detection events **402**, a visibility map **405**, or the like.

[0067] The management system **410** can include a graph control unit **412** to select one or more of the object models implemented by the graph system **420** to utilize in the classification of sensor measurement data **401** associated with one of the detection events **402**. In some embodiments, the graph control unit **412** can select object models based, at least in part, on an association between detection events determined by the event association unit **411**. For example, a previous success or failure in classifying sensor measurement data **401** associated with one detection event can be utilized to select which object models to utilize in attempting to classify sensor measurement data **401** of another detection event. The graph control unit **412** also can select one or more of the object models based on the hypothesis information **403**, processing constraints of hardware implementing the classification system **400**, timing constraints associated with generating the classification **406**, or the like.

[0068] The graph control unit **412** can generate the graph control signaling **414**, which can be configured to direct the graph system **420** to apply selected object models to the matchable representation of the sensor measurement data **401**. The management system **410** can provide the graph control signaling **414** to the graph system **420** along with the matchable representation of the sensor measurement data **401** and optionally the estimated distance between the vehicle and the sensor measurement data **401**.

[0069] In some embodiments, the graph system **420** can include multiple classification graphs **421-1** to **421-N**, each to implement one or more object models describing types of objects capable of being located proximate to the vehicle, such as another vehicle, a pedestrian, a cyclist, an animal, a

static object, or the like. The classification graphs **421-1** to **421**-N can each include one or more nodes having matchable data corresponding to a subset of the various poses, orientations, transitional states, potential deformations, textural features, or the like, in the object model. The management system **410** can select the object models by selecting one or more of the classification graphs **421-1** to **421**-N or selecting one or more of the nodes in the classification graphs **421-1** to **421**-N. Embodiments of a node-based classification graph will be described below in FIG. **5**A.

[0070] FIG. **5**A illustrates an example classification graph **500** in a machine learning object classifier implementation of a classification system according to various examples. Referring to FIG. **5**A, the classification graph **500** can include multiple nodes, each configured to include matchable data corresponding to a subset of the various poses, orientations, transitional states, potential deformations, textural features, inter-frame or inter-scan differences, or the like, in the object model. The nodes can compare a matchable representation of sensor measurement data to the matchable data to determine a match distance. In some embodiments, the matchable representation of sensor measurement data can be cross-checked in the computational nodes. For example, a shape in the matchable representation of sensor measurement data can cross-checked with a skeleton in the matchable representation of sensor measurement data to ascertain a congruency between the shape and the skeleton of the matchable representation. In another example, textural features in the matchable representation of sensor measurement data, such as facial features of a pedestrian, can be cross-checked with the skeleton in the matchable representation of sensor measurement data to ascertain congruency between the textural features and the skeleton of the matchable representation, such whether the facial features are located on an expected portion of the skeleton. The match distance can identify an object type associated with the node, such as another vehicle, a pedestrian, a cyclist, an animal, a static object, or the like, and a confidence level that the matchable representation of the sensor measurement data utilized to generate the match distance corresponds to the object type and an amount of object matching to different object segments. The amount of matching can correspond to matching of object structure, object shape, and object texture in different object segments.

[0071] The classification graph **500** can arrange or organize the nodes in a variety of formats. For example, the classification graph **500** can arrange the nodes so that they are independent of each other. In the embodiments shown in FIG. **5**A, the classification graph **500** can arrange the nodes to be interdependent, for example, to have some nodes corresponding to a coarse-level of detail and other nodes corresponding to finer details for the coarse-level nodes. The classification graph **500** can include an initial node **510** to receive input data **501**, such as a matchable representation of sensor measurement data and an estimated distance of the sensor measurement data from the vehicle. The initial node **510** may include matchable data having a coarse-level of detail corresponding to subsets of the various poses, orientations, transitional states, potential deformations, textural features, or the like, in the object model. For example, when the initial node corresponds to a pedestrian, the matchable data may include to a coarse-level outline of a person with a particular pose, orientation, transitional states, potential deformations of the outline, or the like. In some embodiments, the finer details, such as position of hands, fingers, facial details, or the like, may be omitted from the initial node **510**.

[0072] The classification graph **500** can include delta nodes **521-529**, which can receive the input data **501** from the previously executed node. The delta node **521-529** may include matchable data having a finer-level of detail that the initial node **510** corresponding to subsets of the various poses, orientations, transitional states, potential deformations, textural features, or the like, in the object model. In some embodiments, the delta nodes **521-529** may include matchable data corresponding to differences with the initial node **510**, for example, to avoid duplicative matching of the matchable representation to the matchable data.

[0073] By arranging the nodes in a coarse-level detail to finer-level of detail, the classification system may be able to first determine an object type, for example, that the sensor measurement data corresponds to a person, before determining finer-level details, such as whether the person has a particular orientation, hand position, facial features, or the like. This coarse-level to finer-level detection order can allow the classification system to make the decision when to stop traversing the nodes in the classification graph, reallocate processing resources to different classification activities, control power consumption, or the like. Although FIG. **5**A shows the classification graph **500** including one initial node **510**, in some embodiments, the classification graph **500** can include multiple initial nodes and have multiple sets of delta nodes branching from those initial nodes.

[0074] Referring back to FIG. **4**, each selected node in the graph system **420** can perform a comparison between a portion of the matchable data in the object model to the matchable representation of the sensor measurement data **401**, and generate a match distance **415**, for example, by generating a confidence level that the matchable representation of the sensor measurement data **401** corresponds to the matchable data of the node. The match distance **415** can identify an object type, such as another vehicle, a pedestrian, a cyclist, an animal, a static object, or the like, and a confidence level that the matchable representation of the sensor measurement data **401** utilized to generate the match distance **415** corresponds to the object type. The match distance **415** also can identify granular information about the object type, for example, corresponding to a pose, an orientation, a state, or the like, of the object type. Embodiments of comparing sensor measurement data **401** to matchable data in a node of classification graph will be described below in FIG. **5**B.

[0075] FIG. **5**B illustrates an example flow for comparing sensor measurement data **530** to matchable data **550** in a node of classification graph according to various embodiments. Referring to FIG. **5**B, the sensor measurement data **530** may include a set of data points, such as LIDAR points or point cloud, which a sensor fusion system identified as corresponding to a detection event. As discussed above, the sensor measurement data **530** may be converted into a matchable representation **540**, for example, by determining a boundary outline for the sensor measurement data **530** and creating surface area data within the boundary outline based, at least in part, on the sensor measurement data **530**. The sensor measurement data **530** also may be converted into a matchable representation **540** by skeletonizing the sensor measurement data **530**, for example, by generating a shape

outline or structural body corresponding to the sensor measurement data **530**, and thin the shape outline or structural body into lines. Different points associated with the lines, such as endpoint, vertices where multiple lines connect, or the like, can be identified, and a skeleton of the sensor measurement data **530** can be generated from the lines, the identified points, and relationships between the lines or connectivity of the lines. In some embodiments, the matchable data **550** can correspond to a vehicle classification, which can allow the sensor measurement data **530** to be viewed from a bird's eye view. The sensor measurement data **530** can be skeletonized in the bird's eye view, for example, by generating lines that connect edges of the sensor measurement data **530**. Different points associated with the lines, such as endpoint, vertices where multiple lines connect, or the like, can be identified, and a skeleton of the sensor measurement data **530** can be generated from the lines, the identified points, and relationships between the lines or connectivity of the lines. The skeleton of the sensor measurement data **530** can be included in the matchable representation **540**. The matchable representation **540** can be a two-dimensional representation of the sensor measurement data **530** or a three-dimensional representation of the sensor measurement data **530**.

[0076] The node of the classification graph can compare the matchable representation **540** to matchable data **550** corresponding to a portion of an object model, and generate a match distance to indicate a difference between the matchable representation **540** and the matchable data **550**. In some embodiments, the node can divide the matchable data **550** into different portions or segments, for example, using a fixed-size or fixed grid segmentation, a natural segmentation based on the underlying object described by the object model, or the like, and generate the match distance to indicate differences between the matchable representation **540** and the matchable data **550** for each segment or portion.

[0077] The node can compare the matchable representation **540** to the matchable data **550** in a number of ways. The comparison by the node can determine an overlap between the matchable representation **540** and the matchable data **550**. The node can generate the match distance based, at least in part, on the determined overlap and a significance of the overlapped portions between the matchable representation **540** and the matchable data **550**.

[0078] The node also can locate and handle occlusions, for example, from a visibility map configured to identify where sensor measurement data can be populated into the environmental model based on a view of the sensors. For example, since each sensor has a measurement coordinate range, the visibility map may identify which portions of environmental model correspond to which measurement coordinate ranges. When the autonomous driving system determines an object, such as pedestrian or vehicle, blocks a portion of a normal measurement coordinate range, the autonomous driving system can update the visibility map to identify the occlusion. The node can utilize the occlusions to determine whether the non-overlapped portions were due to a lack of ability for the sensors to measure a portion of the environment surrounding the vehicle or not. The node can generate the match distance based, at least in part, on the occlusion analysis.

[0079] The node can determine whether non-overlapping portions of the matchable representation **540** have a size and shape that would be expected for the object type described by the matchable data **550**. For example, when the matchable data **550** corresponds to a pedestrian, the non-overlapping portions may correspond to a particular positioning of the arms. The node can utilize the matchable data **550** to determine whether the arms in the matchable representation have an expected shape and size, for example, have a natural position and/or normal size relative to the matchable data **550**. The node can generate the match distance based, at least in part, on whether a size and shape the non-overlapping portions of the matchable representation **540** were expected for the object type described by the matchable data **550**.

[0080] The node can attempt to deform the matchable data **550** so that portions of the matchable representation **540** overlap. The node can identify which portions of the matchable data **550** can be deformed by the node and to what degree the node can deform them. The deformation may alter the shape and size of the portion of the matchable data **550** within a set of deformation rules for the matchable data **550**. The node can generate the match distance based, at least in part, on whether deformation allowed there to be an overlap in the matchable representation **540** and the matchable data **550** and/or a distance of the deformation performed to achieve the overlap.

[0081] When the matchable representation **540** corresponds to a detection event associated with at least one previous detection event, the node can determine whether the matchable representation **540** corresponds to an expected state transition compared to a previous state of sensor measurement data for the previous detection event. For example, after a detection event has been previously classified, future detection events associated with that classified detection event can be checked by nodes to determine whether the sensor measurement data in the future detection events corresponds to a state transition from the previous detection event. The node can generate the match distance based, at least in part, on whether the matchable representation **540** match to the matchable data **550** corresponds to an expected state transition.

[0082] Referring back to FIG. **4**, the management system **410** can receive the match distance **415** from each of the initiated nodes in one or more of the classification graphs **421-1** to **421-N**. The management system **410** can include a classification unit **413** to assign a classification **406** to the sensor measurement data **401** associated with the detection events **402** based, at least in part, on at least one of the match distances **415**. For example, when the match distance **415** has a confidence level above a predetermined threshold level, the classification unit **413** can generate the classification **406** for the sensor measurement data **401** associated with the detection events **402**.

[0083] Prior to assigning a classification **406**, the management system **410** also may perform a cross-check of the classification against other sensor data. For example, when the graph system **420** generates a match distance **415** based on LIDAR or RADAR data, the management system **410** or the graph system **420** may analyze image data from a camera sensor to confirm the classification identified in the match distance. The management system **410** or the graph system **420** also may generate a perspective view, such as a bird's eye view or an image view, which can overlay sensor measurement data from multiple different sensor modalities. The management system **410** or the graph system **420** may generate perspective or parallel projections of the sensor

measurement data **401**, such as an orthographic projection onto a ground plane or a perspective projection to the front-facing camera image plane. The management system **410** or the graph system **420** may utilize these generated views or projections to perform cross-checking of the classification indicated from the match distances. In some embodiments, the assignment of the classification **406** or a confidence level associated with the assigned classification **406** also can be based on an assigned classification for any previous detection events associated with the current detection event.

[0084] In some embodiments, the classification unit **413** also can select additional nodes in one or more of the classification graphs **421-1** to **421**-N to initiate based on the match distances **415**, and direct the graph system **420** via the graph control signaling **414**. For example, when the match distance **415** has a confidence level below a predetermined threshold level, the classification unit **413** can determine which of the remaining nodes to initiate or traverse towards in one or more of the classification graphs **421-1** to **421**-N.

[0085] In some embodiments, the traversal through the nodes in the selected classification graphs **421-1** to **421**-N, for example, deciding which nodes should attempt to classify the sensor measurement data **401** and in what order, can be performed independently by the graph system **420**. The classification unit **413** may utilize the match distance **415** to prompt the graph system **420** to cease traversal for different node paths in the classification graphs **421-1** to **421**-N or to direct resources of the graph system **420** towards one or more of the node paths based on the match distances **415**.

[0086] FIG. **6** illustrates an example flowchart for classification of sensor measurement data according to various examples. Referring to FIG. **6**, in a block **601**, a computing system can receive sensor measurement data from different types of sensors in a vehicle. The sensor measurement data can include measurements or raw data from different types or modalities of sensors in a vehicle. For example, the sensor measurement data can include data measured by at least one image capture device, one or more RADAR sensors, one or more LIDAR sensors, one or more ultrasonic sensors, from sources external to the vehicle, such as in a vehicle-to-vehicle (V2V) communication, a vehicle-to-infrastructure (V2I) communication, a vehicle-to-pedestrian (V2P) communication, a vehicle-to-device (V2D) communication, a vehicle-to-grid (V2G) communication, or generally a vehicle-to-everything (V2X) communication, or the like. The computing system can receive the sensor measurement data in frames or scans over time, such as periodically, intermittently, in response to sensing events, or the like. The computing system can spatially align and temporally align the sensor measurement data in an environmental model.

[0087] In a block **602**, the computing system can identify detection events corresponding to possible objects proximate to the vehicle based on the sensor measurement data. The computing system can detect patterns in the sensor measurement data indicative of the possible objects proximate to the vehicle, and output the detected patterns as a detection event. In some embodiments, the computing system can compare detected patterns from multiple different sensor modalities, and identify a detection event from the detected patterns based on a spatial-alignment between the detected patterns, a temporal-alignment between the detected patterns, a state of the data in the between the detected patterns, or the like. The state of the data can correspond to a velocity of the data corresponding to the detected patterns, which can be determined by the sensors and included in the sensor measurement data or determined based on the inter-frame differences. When the computing system utilizes a single sensor modality or sensor type to generate a detection event, the detection event may be called a sensor detection event. When the computing system utilizes associated or correlated patterns from multiple different sensor modalities to generate a detection event, the detection event may be called a fused sensor detection event.

[0088] The computing system also can generate hypothesis information having confidence levels corresponding to properties of the detection events. The hypothesis information can describe various properties or characteristics associated with the detection events and provide confidence levels corresponding to those properties or characteristics. In some embodiments, the properties or characteristics associated with the detection events can include unity, velocity, orientation, center of gravity, existence, size, and novelty of the sensor measurement data corresponding to the detection event. The unity characteristic can identify whether the sensor measurement data corresponds to a single possible object or multiple possible objects proximate to each other, which can help a machine learning classifier select other node or classification graphs corresponding to different portions of an object model. The velocity characteristic can identify at least one velocity associated with the sensor measurement data. The orientation characteristic can identify a directionality of the sensor measurement data and/or an angle associated with the possible object relative to the vehicle. The center of gravity characteristic can identify a center of the possible object or center of a bounding box corresponding to the sensor measurement data based on a density of the data points associated with the detection event. The existence characteristic can identify whether the possible object identified by the detection event is an actual object proximate to the vehicle. The size characteristic can identify or estimate a real size of the possible object associated with the detection event. The novelty characteristic can identify whether the detection event corresponds to a newly detected pattern or corresponding to a previous detection event.

[0089] In a block **603**, the computing system can convert the sensor measurement data corresponding to the detection event into a matchable representation. In some embodiments, the computing system can generate the matchable representation of the sensor measurement data by determining a boundary outline for the sensor measurement data and creating surface area data within the boundary outline based, at least in part, on the sensor measurement data. For example, the computing system can close or fill-in the surface area of the boundary outline by expanding existing data points in the sensor measurement data, connecting existing data points in the sensor measurement data, or the like.

[0090] The matchable representation of the sensor measurement data can be a two-dimensional representation of the sensor measurement data or a three-dimensional representation of the sensor measurement data. When generating the three-dimensional representation of the sensor measurement data, certain volumetric aspects of the matchable representation may not be observable from the sensor measurement data. In these instances, the computing system may estimate the three-dimensional structure associated with the

sensor measurement data based, at least in part, on symmetry in the sensor measurement data, completeness of the sensor measurement data, curvature sensor measurement data, the hypothesis information, or the like.

[0091] In a block **604**, the computing system can select at least one object model to utilize in a classification of the sensor measurement data corresponding to the detection event. In some embodiments, the computing system can select object models based, at least in part, on an association of the detection event and a previously-received detect event. For example, a previous success or failure in classifying sensor measurement data associated with one detection event can be utilized to select which object models to utilize in attempting to classify sensor measurement data of this instant detection event. The computing system also can select one or more of the object models based on the hypothesis information, processing constraints of computing system hardware, timing constraints associated with generating the classification, or the like.

[0092] In a block **605**, the computing system can compare the matchable representation of the sensor measurement data to the selected object model. The object model can include matchable data for a certain object type, which can have various poses, orientations, transitional states, potential deformations for the poses or orientations, textural features, or the like. The computing system can compare the matchable representation of the sensor measurement data to the matchable data from the selection object model, and generate a match distance based on the comparison. In some embodiments, the match distance can correspond to an overlap between the matchable representation and the matchable data, a significance of the overlapped portions, differences between the matchable representation and the matchable data in the non-overlapped portions, whether the matchable data can be deformed to overlap the matchable representation, whether the matchable representation corresponds to an expected state transition for a previously-received detection event, or the like.

[0093] In a block **606**, the computing system can classify the sensor measurement data corresponding to the detection event based, at least in part, on the comparison. The computing system can receive the match distance from the comparison and classify the sensor measurement data based on the match distance. In some embodiments, the computing system can cross-check a classification indicated by the match distance against additional sensor data, such as image data, in an attempt to confirm the classification indicated by the match distance. In some embodiments, the computing system can generate the classification based, at least in part, on an association of the detection event and a previously-classified detect event.

Illustrative Operating Environment

[0094] The execution of various driving automation processes according to embodiments may be implemented using computer-executable software instructions executed by one or more programmable computing devices. Because these embodiments may be implemented using software instructions, the components and operation of a programmable computer system on which various embodiments of the invention may be employed will be described below.

[0095] FIGS. **7** and **8** illustrate an example of a computer system of the type that may be used to implement various embodiments. Referring to FIG. **7**, various examples may be implemented through the execution of software instructions by a computing device **701**, such as a programmable computer. Accordingly, FIG. **7** shows an illustrative example of a computing device **701**. As seen in FIG. **7**, the computing device **701** includes a computing unit **703** with a processing unit **705** and a system memory **707**. The processing unit **705** may be any type of programmable electronic device for executing software instructions, but will conventionally be a microprocessor. The system memory **707** may include both a read-only memory (ROM) **709** and a random access memory (RAM) **711**. As will be appreciated by those of ordinary skill in the art, both the read-only memory (ROM) **709** and the random access memory (RAM) **711** may store software instructions for execution by the processing unit **705**.

[0096] The processing unit **705** and the system memory **707** are connected, either directly or indirectly, through a bus **713** or alternate communication structure, to one or more peripheral devices **717-723**. For example, the processing unit **705** or the system memory **707** may be directly or indirectly connected to one or more additional memory storage devices, such as a hard disk drive **717**, which can be magnetic and/or removable, a removable optical disk drive **719**, and/or a flash memory card. The processing unit **705** and the system memory **707** also may be directly or indirectly connected to one or more input devices **721** and one or more output devices **723**. The input devices **721** may include, for example, a keyboard, a pointing device (such as a mouse, touchpad, stylus, trackball, or joystick), a scanner, a camera, and a microphone. The output devices **723** may include, for example, a monitor display, a printer and speakers. With various examples of the computing device **701**, one or more of the peripheral devices **717-723** may be internally housed with the computing unit **703**. Alternately, one or more of the peripheral devices **717-723** may be external to the housing for the computing unit **703** and connected to the bus **713** through, for example, a Universal Serial Bus (USB) connection.

[0097] With some implementations, the computing unit **703** may be directly or indirectly connected to a network interface **715** for communicating with other devices making up a network. The network interface **715** can translate data and control signals from the computing unit **703** into network messages according to one or more communication protocols, such as the transmission control protocol (TCP) and the Internet protocol (IP). Also, the network interface **715** may employ any suitable connection agent (or combination of agents) for connecting to a network, including, for example, a wireless transceiver, a modem, or an Ethernet connection. Such network interfaces and protocols are well known in the art, and thus will not be discussed here in more detail.

[0098] It should be appreciated that the computing device **701** is illustrated as an example only, and it not intended to be limiting. Various embodiments may be implemented using one or more computing devices that include the components of the computing device **701** illustrated in FIG. **7**, which include only a subset of the components illustrated in FIG. **7**, or which include an alternate combination of components, including components that are not shown in FIG. **7**. For example, various embodiments may be implemented using a multi-processor computer, a plurality of single and/or multiprocessor computers arranged into a network, or some combination of both.

[0099] With some implementations, the processor unit **705** can have more than one processor core. Accordingly, FIG. **8** illustrates an example of a multi-core processor unit **705** that may be employed with various embodiments. As seen in this figure, the processor unit **705** includes a plurality of processor cores **801**A and **801**B. Each processor core **801**A and **801**B includes a computing engine **803**A and **803**B, respectively, and a memory cache **805**A and **805**B, respectively. As known to those of ordinary skill in the art, a computing engine **803**A and **803**B can include logic devices for performing various computing functions, such as fetching software instructions and then performing the actions specified in the fetched instructions. These actions may include, for example, adding, subtracting, multiplying, and comparing numbers, performing logical operations such as AND, OR, NOR and XOR, and retrieving data. Each computing engine **803**A and **803**B may then use its corresponding memory cache **805**A and **805**B, respectively, to quickly store and retrieve data and/or instructions for execution.

[0100] Each processor core **801**A and **801**B is connected to an interconnect **807**. The particular construction of the interconnect **807** may vary depending upon the architecture of the processor unit **705**. With some processor cores **801**A and **801**B, such as the Cell microprocessor created by Sony Corporation, Toshiba Corporation and IBM Corporation, the interconnect **807** may be implemented as an interconnect bus. With other processor units **801**A and **801**B, however, such as the Opteron™ and Athlon™ dual-core processors available from Advanced Micro Devices of Sunnyvale, Calif., the interconnect **807** may be implemented as a system request interface device. In any case, the processor cores **801**A and **801**B communicate through the interconnect **807** with an input/output interface **809** and a memory controller **810**. The input/output interface **809** provides a communication interface between the processor unit **705** and the bus **713**. Similarly, the memory controller **810** controls the exchange of information between the processor unit **705** and the system memory **707**. With some implementations, the processor unit **705** may include additional components, such as a high-level cache memory accessible shared by the processor cores **801**A and **801**B. It also should be appreciated that the description of the computer network illustrated in FIG. **7** and FIG. **8** is provided as an example only, and it not intended to suggest any limitation as to the scope of use or functionality of alternate embodiments.

[0101] The system and apparatus described above may use dedicated processor systems, micro controllers, programmable logic devices, microprocessors, or any combination thereof, to perform some or all of the operations described herein. Some of the operations described above may be implemented in software and other operations may be implemented in hardware. Any of the operations, processes, and/or methods described herein may be performed by an apparatus, a device, and/or a system substantially similar to those as described herein and with reference to the illustrated figures.

[0102] The processing device may execute instructions or "code" stored in a computer-readable memory device. The memory device may store data as well. The processing device may include, but may not be limited to, an analog processor, a digital processor, a microprocessor, a multi-core processor, a processor array, a network processor, or the like. The processing device may be part of an integrated control system or system manager, or may be provided as a portable electronic device configured to interface with a networked system either locally or remotely via wireless transmission.

[0103] The processor memory may be integrated together with the processing device, for example RAM or FLASH memory disposed within an integrated circuit microprocessor or the like. In other examples, the memory device may comprise an independent device, such as an external disk drive, a storage array, a portable FLASH key fob, or the like. The memory and processing device may be operatively coupled together, or in communication with each other, for example by an I/O port, a network connection, or the like, and the processing device may read a file stored on the memory. Associated memory devices may be "read only" by design (ROM) by virtue of permission settings, or not. Other examples of memory devices may include, but may not be limited to, WORM, EPROM, EEPROM, FLASH, NVRAM, OTP, or the like, which may be implemented in solid state semiconductor devices. Other memory devices may comprise moving parts, such as a known rotating disk drive. All such memory devices may be "machine-readable" and may be readable by a processing device.

[0104] Operating instructions or commands may be implemented or embodied in tangible forms of stored computer software (also known as "computer program" or "code"). Programs, or code, may be stored in a digital memory device and may be read by the processing device. "Computer-readable storage medium" (or alternatively, "machine-readable storage medium") may include all of the foregoing types of computer-readable memory devices, as well as new technologies of the future, as long as the memory devices may be capable of storing digital information in the nature of a computer program or other data, at least temporarily, and as long at the stored information may be "read" by an appropriate processing device. The term "computer-readable" may not be limited to the historical usage of "computer" to imply a complete mainframe, mini-computer, desktop or even laptop computer. Rather, "computer-readable" may comprise storage medium that may be readable by a processor, a processing device, or any computing system. Such media may be any available media that may be locally and/or remotely accessible by a computer or a processor, and may include volatile and non-volatile media, and removable and non-removable media, or any combination thereof.

[0105] A program stored in a computer-readable storage medium may comprise a computer program product. For example, a storage medium may be used as a convenient means to store or transport a computer program. For the sake of convenience, the operations may be described as various interconnected or coupled functional blocks or diagrams. However, there may be cases where these functional blocks or diagrams may be equivalently aggregated into a single logic device, program or operation with unclear boundaries.

CONCLUSION

[0106] While the application describes specific examples of carrying out embodiments of the invention, those skilled in the art will appreciate that there are numerous variations and permutations of the above described systems and techniques that fall within the spirit and scope of the invention as set forth in the appended claims. For example, while specific terminology has been employed above to refer to electronic design automation processes, it should be appre-

ciated that various examples of the invention may be implemented using any desired combination of electronic design automation processes.

[0107] One of skill in the art will also recognize that the concepts taught herein can be tailored to a particular application in many other ways. In particular, those skilled in the art will recognize that the illustrated examples are but one of many alternative implementations that will become apparent upon reading this disclosure.

[0108] Although the specification may refer to "an", "one", "another", or "some" example(s) in several locations, this does not necessarily mean that each such reference is to the same example(s), or that the feature only applies to a single example.

1. A method comprising:

generating, by a computing system, a matchable representation of sensor measurement data collected by sensors mounted in a vehicle;

comparing, by the computing system, the matchable representation of the sensor measurement data to an object model describing a type of an object capable of being located proximate to the vehicle; and

classifying, by the computing system, the sensor measurement data as corresponding to the type of the object based, at least in part, on the comparison of the matchable representation of the sensor measurement data to the object model, wherein a control system for the vehicle is configured to control operation of the vehicle based, at least in part, on the classified type of the object for the sensor measurement data.

2. The method of claim **1**, wherein generating the matchable representation of the sensor measurement data further comprises:

extracting lines corresponding to a structure of the sensor measurement data; and

determining relationships and connectivity between the lines, which generates the matchable representation of the sensor measurement data.

3. The method of claim **1**, wherein generating the matchable representation of the sensor measurement data further comprises:

determining a boundary outline for the sensor measurement data; and

creating surface area data within the boundary outline based, at least in part, on the sensor measurement data, which generates the matchable representation of the sensor measurement data.

4. The method of claim **1**, wherein comparing the matchable representation of the sensor measurement data to the object model further comprises:

selecting at least one classification graph from a plurality of classification graphs based, at least in part, on one or more detection events in the sensor measurement data, wherein each classification graph includes one or more computational nodes to implement the object model; and

traversing the computational nodes in the selected classification graph, which compares the matchable representation of the sensor measurement data to different characteristics of the object described in the object model associated with the selected classification graph.

5. The method of claim **4**, wherein each of the computational nodes includes a representation of the object described in the object model with at least one of a pose of the object, a state of the object, an orientation of the object, a textural feature of the object, an inter-frame difference of the object, or one or more deformations for the object.

6. The method of claim **4**, wherein each of the computational nodes in the selected classification graph is configured to generate a match distance between the matchable representation of the sensor measurement data to the representation of the object included in the corresponding computational node, wherein the traversal of the computational nodes in the selected classification graph is based, at least in part, on the match distance generated by one or more of the computational nodes.

7. The method of claim **1**, further comprising estimating, by the computing system, a distance of the object associated with the sensor measurement data from the vehicle, wherein comparing the matchable representation of the sensor measurement data to the object models is based, at least in part, on the estimated distance of the object from the vehicle, a center of gravity of the sensor measurement data, or a center of a bounding box corresponding to the sensor measurement data.

8. An apparatus comprising at least one memory device storing instructions configured to cause one or more processing devices to perform operations comprising:

generating a matchable representation of sensor measurement data collected by sensors mounted in a vehicle;

comparing the matchable representation of the sensor measurement data to an object model describing a type of an object capable of being located proximate to the vehicle; and

classifying the sensor measurement data as corresponding to the type of the object based, at least in part, on the comparison of the matchable representation of the sensor measurement data to the object model, wherein a control system for the vehicle is configured to control operation of the vehicle based, at least in part, on the classified type of the object for the sensor measurement data.

9. The apparatus of claim **8**, wherein generating the matchable representation of the sensor measurement data further comprises:

extracting lines corresponding to a structure of the sensor measurement data; and

determining relationships and connectivity between the lines, which generates the matchable representation of the sensor measurement data.

10. The apparatus of claim **8**, wherein generating the matchable representation of the sensor measurement data further comprises:

determining a boundary outline for the sensor measurement data; and

creating surface area data within the boundary outline based, at least in part, on the sensor measurement data, which generates the matchable representation of the sensor measurement data.

11. The apparatus of claim **8**, wherein comparing the matchable representation of the sensor measurement data to the object model further comprises:

selecting at least one classification graph from a plurality of classification graphs based, at least in part, on one or more detection events in the sensor measurement data, wherein each classification graph includes one or more computational nodes to implement the object model; and

traversing the computational nodes in the selected classification graph, which compares the matchable representation of the sensor measurement data to different characteristics of the object described in the object model associated with the selected classification graph.

12. The apparatus of claim **11**, wherein each of the computational nodes includes a representation of the object described in the object model with at least one of a pose of the object, a state of the object, an orientation of the object, a textural feature of the object, an inter-frame difference of the object, or one or more deformations for the object.

13. The apparatus of claim **11**, wherein each of the computational nodes in the selected classification graph is configured to generate a match distance between the matchable representation of the sensor measurement data to the representation of the object included in the corresponding computational node, wherein the traversal of the computational nodes in the selected classification graph is based, at least in part, on the match distance generated by one or more of the computational nodes.

14. The apparatus of claim **8**, wherein the instructions are further configured to cause the one or more processing devices to perform operations comprising estimating a distance of the object associated with the sensor measurement data from the vehicle, wherein comparing the matchable representation of the sensor measurement data to the object models is based, at least in part, on the estimated distance of the object from the vehicle, a center of gravity of the sensor measurement data, or a center of a bounding box corresponding to the sensor measurement data.

15. A system comprising:

a memory device configured to store machine-readable instructions; and

a computing system including one or more processing devices, in response to executing the machine-readable instructions, configured to:

generate a matchable representation of sensor measurement data collected by sensors mounted in a vehicle;

compare the matchable representation of the sensor measurement data to an object model describing a type of an object capable of being located proximate to the vehicle; and

classify the sensor measurement data as corresponding to the type of the object based, at least in part, on the comparison of the matchable representation of the sensor measurement data to the object model, wherein a control system for the vehicle is configured to control operation of the vehicle based, at least in part, on the classified type of the object for the sensor measurement data.

16. The system of claim **15**, wherein the one or more processing devices, in response to executing the machine-readable instructions, are configured to:

extract lines corresponding to a structure of the sensor measurement data; and

determine relationships and connectivity between the lines, which generates the matchable representation of the sensor measurement data.

17. The system of claim **15**, wherein the one or more processing devices, in response to executing the machine-readable instructions, are configured to:

determine a boundary outline for the sensor measurement data; and

create surface area data within the boundary outline based, at least in part, on the sensor measurement data, which generates the matchable representation of the sensor measurement data.

18. The system of claim **15**, wherein the one or more processing devices, in response to executing the machine-readable instructions, are configured to:

select at least one classification graph from a plurality of classification graphs based, at least in part, on one or more detection events in the sensor measurement data, wherein each classification graph includes one or more computational nodes to implement the object model; and

traverse the computational nodes in the selected classification graph, which compares the matchable representation of the sensor measurement data to different characteristics of the object described in the object model associated with the selected classification graph.

19. The system of claim **18**, wherein each of the computational nodes in the selected classification graph is configured to generate a match distance between the matchable representation of the sensor measurement data to the representation of the object included in the corresponding computational node, wherein the traversal of the computational nodes in the selected classification graph is based, at least in part, on the match distance generated by one or more of the computational nodes.

20. The system of claim **15**, wherein the one or more processing devices, in response to executing the machine-readable instructions, are configured to estimate a distance of the object associated with the sensor measurement data from the vehicle, and compare the matchable representation of the sensor measurement data to the object models based, at least in part, on the estimated distance of the object from the vehicle, a center of gravity of the sensor measurement data, or a center of a bounding box corresponding to the sensor measurement data.

* * * * *