



(43) International Publication Date
06 September 2019 (06.09.2019)

(51) International Patent Classification:
G08G 1/00 (2006.01)

(21) International Application Number:
PCT/US2019/019890

(22) International Filing Date:
27 February 2019 (27.02.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/635,701 27 February 2018 (27.02.2018) US
62/729,350 10 September 2018 (10.09.2018) US

(71) Applicant: NAUTO, INC. [US/US]; 220 Portage Avenue, Palo Alto, CA 94306 (US).

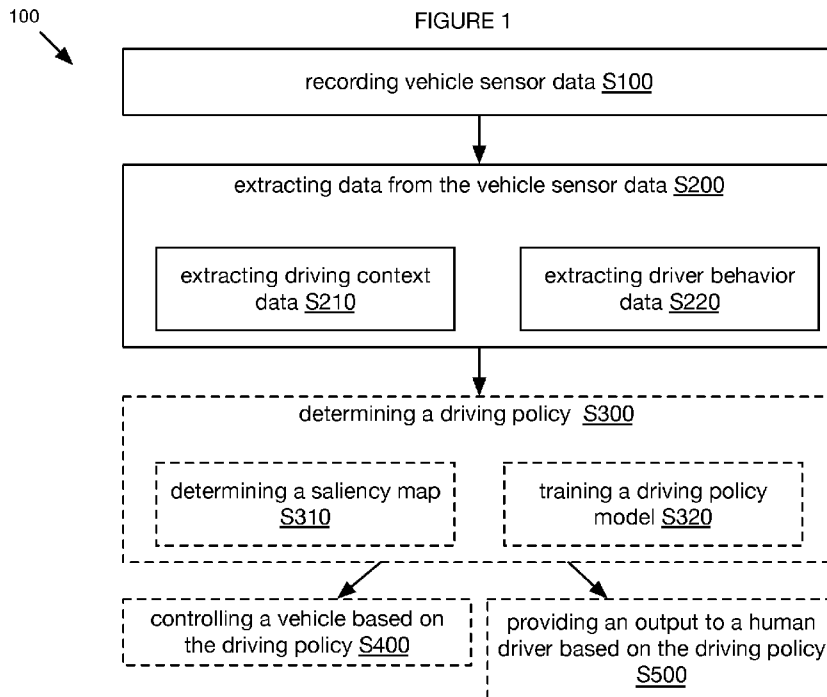
(72) Inventors: SATZODA, Ravi Kumar; 220 Portage Avenue, Palo Alto, CA 94306 (US). SATHYNARAYANA, Suchitra; 220 Portage Avenue, Palo Alto, CA 94306 (US). LEVKOVA, Ludmilla; 220 Portage Avenue, Palo Alto, CA 94306 (US). HECK, Stefan; 220 Portage Avenue, Palo Alto, CA 94306 (US).

(74) Agent: LIN, Diana; 500 3rd Street, Suite 215, San Francisco, CA 94107 (US).

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

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

(54) Title: METHOD FOR DETERMINING DRIVING POLICY



(57) Abstract: Systems and methods for driving. A driving data set for each of plurality of human-driven vehicles is determined. For each driving data set, exterior scene features of an exterior scene of the respective vehicle are extracted from the exterior image data. A driving response model is trained based on the exterior scene features and the vehicle control inputs from the selected driving data sets.



WO 2019/169031 A1

Published:

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

METHOD FOR DETERMINING DRIVING POLICY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of US Provisional Application number 62/635,701 filed 27-FEB-2018, and US Provisional Application number 62/729,350 filed 10-SEP-2018, which are incorporated in their entireties by this reference.

TECHNICAL FIELD

[0002] This invention relates generally to the automotive vehicle control field, and more specifically to a new and useful method for determining driving policy in the automotive vehicle control field.

BACKGROUND

[0003] Automotive accidents are a major cause of deaths and injuries to human drivers. In order to improve safety and significantly reduce the number of fatalities, autonomous driving systems and control methods are being considered as an effective solution. Machine learning can play a significant role in developing such autonomous driving systems and control methods, wherein computing systems can be trained to drive safely and with minimal intervention from human drivers, according to a set of driving behavior rules for various real-world situations that can be collectively defined as driving policy. However, training such systems can require large quantities of accurate and salient data, and data saliency can be difficult to determine without excessive time and expense (e.g., through the use of human labeling, filtering, and/or other manual techniques for determination of data saliency, etc.).

[0004] Thus, there is a need in the automotive field to create a new and useful method for determining driving policy. This invention provides such a new and useful method.

BRIEF DESCRIPTION OF THE FIGURES

- [0005]** FIGURE 1 depicts a flowchart of a variation of the method;
- [0006]** FIGURE 2 depicts a schematic diagram of an onboard vehicle system that can be used to implement portions of variations of the method for determining driving policy;
- [0007]** FIGURE 3 depicts a schematic diagram of a saliency map associated with an example implementation of the method for determining driving policy;
- [0008]** FIGURE 4 depicts a schematic of a portion of an example implementation of the method for determining driving policy;
- [0009]** FIGURES 5A-B depict flowchart diagrams of variations of a method; and
- [0010]** FIGURES 6A and 6B depict specific examples of model training and use, respectively.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

- [0011]** The following description of the preferred embodiments of the invention is not intended to limit the invention to these preferred embodiments, but rather to enable any person skilled in the art to make and use this invention.
- [0012]** As shown in FIGURES 1 and 5, the method 100 for determining driving policy includes: recording vehicle sensor data at an onboard vehicle system (e.g., 510, shown in FIGURE 2) during a vehicle event S100; extracting vehicle event data and driver behavior data from the vehicle sensor data S200; and, determining a driving policy based on the vehicle event data in combination with the driver behavior data S300. The method 100 can optionally include controlling a vehicle (e.g., 501, shown in FIGURE 2) based on the driving policy S400, providing output to a human driver based on the driving policy S500, and/or any other suitable blocks or processes.
- [0013]** The method 100 functions to correlate driver behavior with aspects of vehicle events (e.g., by determining the relative saliency of various portions of vehicle

events), and to determine driving policy rules based on this correlation that enable vehicle control systems to emulate and/or improve upon the positive aspects of the driving behavior. The method 100 can also function to develop driving policy rules that improve upon the negative aspects of human driving behavior (e.g., human loss of focus or attention, comparatively slow human cognition and/or perception speed, etc.). The method 100 can also function to train models (e.g., driving policy models, inference models, decision making models, etc.) using correlated driver behavior data and vehicle event data (e.g., in the form of a saliency map of the vehicle event at each time point during the vehicle event). In variants, the models can be used to control autonomous or semi-autonomous vehicles, particularly in complex driving environments, such as intersections. The models can make better (e.g., make safer, more efficient, more predictable, etc.) decisions than conventional models, since the models were developed on real-world data collected in similar complex driving environments. In specific examples, the method can leverage human driver behavior during specific driving events to generate data (e.g., labeled data, supervised training data) to train inference systems with real-world naturalistic driving scenarios, such that the resultant models can behave (e.g., react, control vehicles, etc.) similar to or better than human drivers. The method 100 can also function to generate a training dataset (e.g., from saliency-mapped vehicle event data, from vehicle event data labeled using a set of driving policy rules, determined via one or more variations of a portion of the method, etc.) that can be utilized (e.g., by a third party, by an autonomous vehicle system, etc.) for training vehicle control models. The method 100 can also function to control a vehicle (e.g., automatically control an autonomous vehicle) according to a driving policy determined in accordance with one or more variations of a portion of the method. The method 100 can also function to estimate what another human-driven vehicle will do, and can feed said output to a secondary autonomous vehicle control model. The method 100 can also function to improve the performance of a human driver using the output of a driving policy model (e.g., developed based on a population of human drivers, developed based on historical consideration of one or more human drivers over time, etc.), such as by providing the output to the human

driver (e.g., real-time coaching via audiovisual stimuli, after-the-fact coaching via summary performance reports, etc.). However, the method 100 can additionally or alternatively have any other suitable function.

1. Benefits

[0014] Variations of the technology can afford several benefits and/or advantages.

[0015] First, variations of the technology can enable skilled driving behaviors to be identified, recorded, and utilized to generate automated or semi-automated vehicle control systems exhibiting equivalent and/or superior driving skills.

[0016] Second, variations of the technology can enable unskilled drivers to improve their skills through coaching based on driving policy determined based on skilled drivers (e.g., without direct interaction between skilled and unskilled drivers, such as through an in-person driver training program).

[0017] Third, variations of the technology can enable the training and/or evaluation of computational models for vehicle control according to determined driving policies. For example, the method can include filtering a dataset for vehicle event data associated with skilled driving (e.g., by labeling vehicle event data using a driving policy model generated in accordance with a variation of a portion of the method, using a second scoring model, etc.), and using the filtered dataset to train a vehicle control model to incorporate the methodologies of skilled driving. In other examples, the method can include comparing the output of a vehicle control model to a filtered dataset of vehicle event data associated with skilled driving, and/or to a driving policy model generated using such a filtered dataset, to evaluate the output of the vehicle control model (e.g., use the collected data to test whether a vehicle control model is satisfactory or would make satisfactory vehicle control decisions in complex driving environments). In another example, the collected data can be used to train a model that dictates when and/or whether an autonomous vehicle control system should be disengaged (e.g., when and/or whether a human driver should regain control of the vehicle).

[0018] Fourth, variations of the technology can confer improvements in computer-related technology (e.g., vehicle telematics, computational modeling associated with vehicle movement characteristics, etc.) by leveraging non-generic vehicle event data (e.g., extracted from exterior image data, extracted from correlated interior-exterior data, etc.), driver behavior data (e.g., extracted from interior image data, extracted from correlated interior-exterior data, etc.), and/or other suitable data from one or more devices (e.g., non-generalized onboard vehicle systems), sensor systems associated with the vehicle and/or surroundings of the vehicle, and any other suitable systems to improve accuracy of driving policy determination related to vehicle operation and/or vehicle movement characteristics (e.g., which can thereby enable appropriately generated and/or timed user-related actions, vehicle control instructions, etc.). In examples, the technology can confer improvements in the application of such technology by enabling convenient, unobtrusive, accurate, and/or skillful autonomous or semi-autonomous vehicle control matching or exceeding the performance of skilled human drivers, as well as improved vehicle control over time.

[0019] Fifth, variations of the technology can provide technical solutions necessarily rooted in computer technology (e.g., utilizing different computational models to determine driving policy based on data streams from sensor systems, etc.) to overcome issues specifically arising with computer technology (e.g., issues surrounding how to leverage correlated interior-exterior image data in association with vehicle events; issues surrounding accurately and appropriately performing control actions for different vehicle events, vehicle event types, and the like; etc.). In another example, the technology can apply computer-implemented rules (e.g., feature engineering rules for processing sensor data into an operable form for generating features; sensor data collection and/or processing rules for data from onboard vehicle systems and/or associated computing devices, mobile devices, sensor systems; etc.).

[0020] Sixth, variations of the technology can confer improvements in the functioning of computational systems themselves. For example, the technology can improve upon the processing of collected non-generic data (e.g., by filtering the collected

sensor data based on the saliency of the data, enabling the most salient data to be focused upon and processed and the least salient data to be ignored or de-weighted during processing).

[0021] Seventh, by collecting training data from real, human-controlled driving sessions, the method can collect naturalistic driving responses in real-world driving contexts. Training autonomous vehicle control models to emulate naturalistic driving responses can be particularly useful in hybrid driving environments where autonomous vehicles share the road with human-driven vehicles, since the human drivers may expect the autonomous vehicles to have human-like responses to driving events.

[0022] Eighth, by collecting said data from a plurality of drivers, vehicles, and/or driving sessions, the method can collect data for edge-case driving events (e.g., rare driving events, difficult-to-simulate events, etc.) and/or complex driving environments.

[0023] Ninth, by collecting interior data in addition to exterior data, the method can determine the human driver's gaze (e.g., from the interior data) relative to the external scene (e.g., from the exterior data), and determine a region of interest. This region of interest can be used to determine which portion of the external scene to pay attention to (e.g., wherein the region(s) of interest can be used to train an attention model or scanning model that subsequently feeds in to the driving policy model), which can function to reduce the processing resources required to run the driving policy model.

[0024] However, variations of the method can offer any other suitable benefits and/or advantages.

2. System

[0025] The method can be performed at least in part by a sensing and computing system on-board the vehicle (e.g., an onboard vehicle system, e.g., 510), but can additionally or alternatively be performed at least in part by a remote computing system (e.g., 520 shown in FIGURE 2), such as a server system, a user device (e.g., a smartphone, a tablet, etc.), or by any other suitable set or network of computing systems. The method is preferably performed using data sampled by the onboard vehicle system (e.g., vehicle

sensor data), but can additionally or alternatively be performed using auxiliary vehicle data (e.g., signals sampled by the other vehicle sensors besides those of the onboard vehicle system, vehicle data retrieved from a database, intrinsic vehicle data associated with the vehicle itself and stored at the onboard vehicle system, etc.), other vehicles' data (e.g., received from the source vehicle, a database, or any other suitable remote computing system), aggregate population data, historic data (e.g., for the vehicle, driver, geographic location, etc.), or any other suitable data from any other suitable source. The onboard vehicle system (e.g., 510) can function to capture, record, or otherwise suitably obtain vehicle sensor data corresponding to the vehicle surroundings during a vehicle event (e.g., the event scene, driving scene, etc.) while simultaneously capturing, recording, or otherwise suitably vehicle sensor data corresponding to the driver (e.g., for use in determining the driver behavior) during a vehicle event. However, the onboard vehicle system can otherwise suitably capture correlated interior-exterior data usable to determine the driving policy of the driver.

[0026] The onboard vehicle system (e.g., 510) can include a processing system (e.g., a set of GPUs, CPUs, microprocessors, TPUs, vehicle computing systems, etc.), storage system (e.g., RAM, Flash), communication system, sensor set (e.g., 531-533 shown in FIGURE 2), power system (e.g., battery, vehicle power connector, photovoltaic system, etc.), CAN bus interface (e.g., wired or wireless), housing, or any other suitable component. The communication system can include telemetry systems (e.g., for vehicle-to-vehicle, vehicle-to-infrastructure, vehicle-to-remote computing system, or other communications), wireless systems (e.g., cellular, WiFi or other 802.11x protocols, Bluetooth, RF, NFC, etc.), wired systems (e.g., Ethernet, vehicle bus connections, etc.), or any other suitable communication systems. The sensors (e.g., 531-534) shown in FIGURE 2) can include: cameras (e.g., wide angle, narrow angle, or having any other suitable field of view; visible range, invisible range, IR, multispectral, hyperspectral, or sensitive along any suitable wavelength; monocular, stereoscopic, or having any suitable number of sensors or cameras; etc.), kinematic sensors (e.g., accelerometers, IMUs, gyroscopes, etc.), optical systems (e.g., ambient light sensors), acoustic systems (e.g., microphones,

speakers, etc.), range-finding systems (e.g., radar, sonar, TOF systems, LIDAR systems, etc.), location systems (e.g., GPS, cellular trilateration systems, short-range localization systems, dead-reckoning systems, etc.), temperature sensors, pressure sensors, proximity sensors (e.g., range-finding systems, short-range radios, etc.), or any other suitable set of sensors.

[0027] In one variation, an example of which is shown in FIGURE 2, the onboard vehicle system 510 at which at least a portion of the method 100 is implemented includes a set of internal sensors (e.g., 531), a set of exterior sensors (e.g., 532), and a processing system. The internal sensors (e.g., internal-facing camera 535, microphones, etc.) can be directed toward and monitor the vehicle interior, more preferably the driver volume (e.g., the volume of the interior in which a vehicle driver is and/or would be situated during driving of the vehicle) but alternatively or additionally any suitable interior volume. The exterior sensors (e.g., exterior-facing camera 536, rangefinding sensors, etc.) are preferably directed outward from the vehicle, and preferably include a region in front of the vehicle (e.g., region preceding the vehicle along the vehicle trajectory, region proximal the driving volume and encompassing the vehicle drivetrain longitudinal vector, etc.), but can alternatively be directed toward the vehicle side(s), top, bottom, rear, or any other suitable region exterior the vehicle and/or including the vehicle surroundings. The sensors are preferably statically mounted to the vehicle 501 and/or each other (e.g., via the housing), but can be movably mounted by a gimbal, damping system, or other motion mechanism.

[0028] Each camera's intrinsic parameters are preferably known (e.g., wherein the processing system processing the camera images can store an intrinsic matrix for each camera), but can alternatively be unknown and/or calibrated on-the-fly. The extrinsic parameters relating the internal-facing camera (e.g., included in 531) with the external-facing camera (e.g., included in 532) is preferably also known (e.g., wherein the processing system processing the respective camera images stores an extrinsic matrix for the sensor system), but can alternatively be unknown and/or calibrated on-the-fly. The intrinsic and extrinsic matrices are preferably held constant (e.g., wherein the camera

components are assumed to not warp or shift, and the interior-facing camera and the exterior-facing camera are assumed to remain statically coupled by the housing), but can alternatively be dynamically determined or otherwise determined. In one example, a portion of the interior images can be pre-associated with a portion of the exterior images, wherein the mapping can be dynamically determined based on the extrinsic matrix, predetermined (e.g., during calibration), or otherwise determined. The interior-facing camera and exterior-facing cameras are preferably synchronized in time (e.g., by sharing a common clock, calibrating against an external temporal reference, such as a GPS clock, etc.), but the resultant images can be otherwise associated with each other.

[0029] In one example, the system can include or interact with an OBD II scanner communicably connected to the onboard vehicle system (e.g., wirelessly, via a wired connection). In a second example, the vehicle ECU(s) can directly communicate with the onboard vehicle system. However, the onboard vehicle system can receive information from the vehicle control system in any other suitable manner.

[0030] In variants in which the resultant models (e.g., driving policy models, attention models, scanning models, etc.) are used to control an autonomous vehicle (or semi-autonomous vehicle), the autonomous vehicle preferably includes external sensors (e.g., distance sensors, rangefinding sensors such as LIDAR, cameras, radar, proximity sensors, etc.) and control inputs (e.g., acceleration, braking, steering, etc.), but can additionally or alternatively include interior sensors or any other suitable set of sensors.

[0031] In some variations, the onboard vehicle system 510 (and/or autonomous vehicle using the trained model(s)) includes a vehicle control subsystem. In some variations, the onboard vehicle system 510 is communicatively coupled to a vehicle control subsystem (e.g., 512 shown in FIGURE 2) that is included in a separate housing from a housing that includes the onboard vehicle system 510. In some variations, the vehicle control subsystem functions to receive control inputs (e.g., control instructions for the control inputs, target control input values, etc.) and control at least one of acceleration, braking, and steering of the vehicle 501 based on the received control inputs.

In some variations, the onboard vehicle system 510 is communicatively coupled to the vehicle control system 512 via either a bus or a local network of the vehicle 501.

3. Method

[0032] As shown in FIGURE 1, the method 100 includes: recording vehicle sensor data at an onboard vehicle system S100; extracting driving context data and driver behavior data from the vehicle sensor data S200; and determining a driving policy based on the driving context data in combination with the driver behavior data S300. The method 100 can optionally include controlling a vehicle based on the driving policy S400; providing output to a human driver based on the driving policy S500; and/or any other suitable blocks or processes.

[0033] The method 100 can be performed (e.g., executed, implemented, etc.) in real- or near-real time, but all or portions of the method can alternatively be performed asynchronously or at any other suitable time. The method is preferably iteratively performed at a predetermined frequency (e.g., every millisecond, at a sampling frequency, etc.), but can alternatively be performed in response to occurrence of a trigger event (e.g., change in the vehicle attitude, change in user distraction levels, receipt of driving session information, receipt of new sensor information, physical vehicle entry into a geographic region associated with high collision risk, object proximity detection, detection of an onset or end of a driving session, etc.), be performed a single time for a driving session, be performed a single time for the vehicle, or be performed at any other suitable frequency.

[0034] One or more variations of the method 100 can be performed for each of a plurality of vehicles, such as vehicles equipped with an onboard vehicle system as described herein (e.g., 510, shown in FIGURE 2), and can be performed for a plurality of driving sessions and/or drivers, thereby generating data sets across multiple vehicles, drivers, and/or driving sessions.

[0035] Block S100 includes recording vehicle sensor data. In some variations, the vehicle sensor data is recorded during a driving session. Block S100 functions to obtain data indicative of the surroundings of a vehicle and the actions of the driver in relation to

the surroundings during a driving-related scenario (e.g., a vehicle event, driving context). The vehicle sensor data is preferably recorded using an onboard vehicle system (e.g., 510) as described above; however, vehicle sensor data can additionally or alternatively be recorded using any suitable sensor system, integrated with and/or distinct from the vehicle (e.g., 501) itself (e.g., the host vehicle, the ego-vehicle, etc.). Vehicle sensor data is thus preferably indicative of the surroundings of a host vehicle and of the interior of the host vehicle (e.g., 501). The collected vehicle sensor data can be associated with: one or more driving contexts, a driver identifier, a driving session, and/or any other suitable information.

[0036] Block S100 functions to record vehicle sensor data that can be used to generate a driving data set for each of a plurality of human-driven vehicles. In some variations, each driving data set includes sensor data for at least one driving session or driving event of a vehicle. In some variations, each driving data set includes sensor data for at least one maneuver of a driving session. In some variations, at least one maneuver is associated with information indicating a skill metric. In some variations, each driving data set includes sensor information for determining at least one of: a driver ID for each driving data session, a driver attentiveness score for each driving session, a skill metric (e.g., for the driver, for a maneuver), a driver attentiveness score for each driving event represented by the driving data set, and/or any other suitable upstream analysis.

[0037] In some variations, driving data sets can be tagged one or more of: driving event data (e.g., data indicating a detected event associated with the driving data set), data indicating a driving maneuver performed by the human driver in response to an event, driver ID of the driver, the driver control inputs (e.g., acceleration, braking, steering, signaling, etc.), and/or any other suitable data. The driver control inputs can be the vehicle control inputs applied by the driver: simultaneously with driving data set sampling (e.g., encompass the same timeframe as or be within the timeframe of the driving data set); contemporaneous with driving data set sampling (e.g., encompass a timeframe overlapping or encompassing the driving data set timeframe); within a predetermined time window of driving data set sampling (e.g., a predetermined time

window after the driving data set timeframe, such as the next 10 seconds, next 30 seconds, next minute, the next 5 minutes, the next 10 minutes, the time window between 10 seconds to 5 minutes after the driving data set timeframe, etc.); or be the control inputs applied by the driver at any other suitable time relative to the driving data set timeframe. The driving data sets can be tagged or be associated with the data by: the onboard vehicle system 510, the remote computing system 520), and/or any other suitable system.

[0038] In some variations, the vehicle sensor data is recorded during a vehicle event. In some variations, the vehicle sensor data is continuously recorded. In some variations, the vehicle sensor data is discontinuously recorded at periodic or irregular sampling intervals.

[0039] Vehicle sensor data collected in accordance with Block S100 can include synchronous data (e.g., temporally synchronous, spatially synchronized or correlated, etc.) captured from at least two cameras: a first camera (e.g., 536, shown in FIGURE 2) oriented to image outside the vehicle, and a second camera (e.g., 535, shown in FIGURE 2) oriented to image within the vehicle. The vehicle sensor data can additionally or alternatively include location data (e.g., GPS data), motion data (e.g., inertial measurement unit / IMU data), and any other suitable type of sensor data. The synchronized sensor data (e.g., from inside and outside cameras) can be used to correlate driver activities (e.g., driver behavior) to events that are happening outside the vehicle (e.g., vehicle events, diving scenarios, etc.). Vehicle sensor data that is collectively aggregated from one or more data streams preferably includes two-way video data (e.g., inward facing video camera data and outward facing video camera data), and can also include inertial data, gyroscope data, location data, routing data, kinematic data, and other suitable vehicle telemetry data (e.g., collected from an OBD II port via a suitable data connection). However, vehicle sensor data can include any other suitable data.

[0040] In some variations, Block S100 includes sampling synchronized interior sensor data and exterior sensor data for inclusion in a driving data set, as described herein, that also includes vehicle control inputs (e.g., acceleration, steering, braking,

signaling, etc.) associated with the synchronized interior sensor data and exterior sensor data.

[0041] In some variations, block S100 includes detecting one or more predetermined driving events at a vehicle, and sampling the synchronized interior sensor data and exterior sensor data (as described herein) after detecting at least one predetermined driving event. Driving events can include vehicle arrival at an intersection, the vehicle being tailgated by another vehicle, the vehicle tailgating another vehicle, traffic, the vehicle being cut-off by another driver, and the like.

[0042] In some variations a single sensor sampling is performed in response to detection of a driving event. In some variations, several sensor samplings are performed in response to detection of a driving event (e.g., continuous or discontinuous sampling within a predetermined time period or until a stopping condition is satisfied). In some variations, interior sensor data and exterior sensor data are both image data, and at least one predetermined driving event is detected based on sensor data other than the image data of the vehicle (auxiliary sensor data). Auxiliary sensor data can include data generated by kinematic sensors (e.g., accelerometers, IMUs, gyroscopes, etc.), optical systems (e.g., ambient light sensors), acoustic systems (e.g., microphones, speakers, etc.), range-finding systems (e.g., radar, sonar, TOF systems, LIDAR systems, etc.), location systems (e.g., GPS, cellular trilateration systems, short-range localization systems, dead-reckoning systems, etc.), temperature sensors, pressure sensors, proximity sensors (e.g., range-finding systems, short-range radios, etc.), or any other suitable set of sensors.

[0043] In some variations, the interior sensor data includes image data captured by an interior camera (e.g., 535) oriented to image the vehicle interior. In a first variation, the interior image data included in the driving data set include complete frames of captured interior image data. In a second variation, the interior image data included in the driving data set include cropped frames of captured interior image data. For example, a driver face can be identified in the frames of the interior image data, the frames of the interior image data can be cropped to the identified driver face, and the cropped frames can be included in the driving data set instead of the full frames, such a size of the driving

data set can be reduced as compared to a driving data set that includes the full (uncropped) interior image data. In a first example, the cropped frames can be used to determine driving context (e.g., an identification of a current driver, presence of a human driver). In a second example, the cropped frames can be used to determine driver behavior (e.g., gaze, head pose, attentiveness, etc.) of a current driver.

[0044] In some variations, the exterior sensor data includes image data captured by an exterior camera (e.g., 536) oriented to image outside the vehicle. In some variations, the exterior sensor data includes LIDAR data captured by a LIDAR systems oriented to a scene outside the vehicle. In some variations, the exterior sensor data includes a point cloud dataset representing a scene outside the vehicle as sensed by a LIDAR system.

[0045] In some variations, the external scene representation (extracted from the exterior sensor data) can be converted to the output format for a secondary sensor suite (e.g., using a translation module, such as a depth map-to-point cloud converter; etc.). The secondary sensor suite is preferably that of the autonomous vehicle using the trained model(s), but can be any other suitable set of sensors. This translation is preferably performed before external scene feature extraction and/or model training, such that the trained model will be able to accept features from the secondary sensor suite and is independent from the onboard vehicle system's sensor suite and/or features extracted therefrom. However, the translation can be performed at any suitable time, or not performed at all. In some examples, block S100 includes generating a LIDAR point cloud dataset representing a scene outside the vehicle from image data captured by an exterior camera oriented to image outside the vehicle.

[0046] The method can optionally include determining the driving context associated with a set of vehicle sensor data.

[0047] The driving context can be used in multiple ways. In one variation, the vehicle sensor data is collected upon occurrence of a predetermined driving context (e.g., the current driving context satisfying a predetermined set of conditions). This can function to minimize the amount of data that needs to be stored on-board the vehicle

and/or the amount of data that needs to be analyzed and/or transmitted to the analysis system. The driving policy model trained using such data can be specific to the predetermined driving context, a set thereof, or generic to multiple driving contexts. Examples of predetermined driving contexts include: vehicle proximity to complex driving locations, such as intersections (e.g., wherein the vehicle is within a geofence associated with an intersection, when the external sensor measurements indicate an intersection, etc.); vehicle events; autonomous control model outputs having a confidence level lower than a threshold confidence; complex driving conditions (e.g., rain detected within the external image or by the vehicle's moisture sensors); or any other suitable driving context.

[0048] In a second variation, the driving context (e.g., driving context features) is used as the training input, wherein the driver's control inputs (e.g., concurrent with the driving context or subsequent the driving context, within a predetermined timeframe, etc.) are used as the data label. However, the driving context can be otherwise used.

[0049] Driving context can include: driving event(s), location (e.g., geolocation), time, the driving environment (e.g., external scene, including the position and/or orientation of external objects relative to the host vehicle and/or estimated object trajectories; ambient environment parameters, such as lighting and weather, etc.), vehicle kinematics (e.g., trajectory, velocity, acceleration, etc.), next driving maneuver, urgency, or any other suitable driving parameter. The driving context can be determined: in real-time, during the driving session; asynchronously from the driving session; or at any suitable time. The driving context can be determined using: the onboard vehicle system, a remote computing system, and/or any other suitable system. The driving context can be determined based on: the vehicle sensor data, vehicle control data, navigation data, data determined from a remote database, or any other suitable data.

[0050] A vehicle event can include any driving-related, traffic-related, roadway-related, and/or traffic-adjacent event that occurs during vehicle operation. For example, a vehicle event can include an interaction between the ego-vehicle (e.g., the host vehicle, the vehicle on which the onboard vehicle system is located, etc.) and another vehicle (e.g.,

a secondary vehicle), pedestrian, and/or other static or non-static (e.g., moving) object. An interaction can be a collision, a near-collision, an effect upon the driver of the presence of the secondary vehicle or traffic object (e.g., causing the driver to slow down, to abstain from accelerating, to maintain speed, to accelerate, to brake, etc.), typical driving, arrival at a predetermined location or location class (e.g., location within or proximal to an intersection), and/or any other suitable type of interaction. The vehicle event can include a driving maneuver, performed in relation to the ego-vehicle (e.g., by a driver of the ego-vehicle) and/or a secondary vehicle (e.g., by a driver or operator of the secondary vehicle). A driving maneuver can be any operation performable by the vehicle (e.g., a left turn, a right turn, a lane change, a swerve, a hard brake, a soft brake, maintaining speed, maintaining distance from a leading vehicle, perpendicular parking, parallel parking, pulling out of a parking spot, entering a highway, exiting a highway, operating in stop-and-go traffic, standard operation, non-standard operation, emergency action, nominal action, etc.).

[0051] A vehicle event can be of any suitable duration; for example, a vehicle event can be defined over a time period of a driving maneuver, over a time period of a set of related driving maneuvers (e.g., changing lanes in combination with exiting a highway, turning into a parking lot in combination with parking a vehicle, etc.), over a time period encompassing a driving session (e.g., the time between activation of a vehicle and deactivation of the vehicle), continuously during at least a portion of a driving session, of a variable duration based on event characteristics (e.g., over a time period of highway driving that is delimited in real time or after the fact based on recognition of the vehicle entering and/or exiting the highway region), and any other suitable duration or time period associated with a driving session.

[0052] A vehicle event can be determined in real time (e.g., during a driving session made up of a plurality of vehicle events) based on collected vehicle sensor data, subsequent to sensor data collection (e.g., wherein data is recorded, sampled, or otherwise obtained in accordance with one or more variations of Block S100) as at least a

portion of the vehicle event data extraction of Block S200, and/or otherwise suitably determined.

[0053] Vehicle event (driving event) detection can be performed by a model, such as an artificial neural network (e.g., a convolutional neural network), Bayesian model, a deterministic model, a stochastic and/or probabilistic model, a rule-based model, and any other suitable model. Driving event detection is preferably performed, at least in part, onboard the vehicle (e.g., at an onboard vehicle system, a vehicle computing unit, an electronic control unit, a processor of the onboard vehicle system, a mobile device onboard the vehicle, etc.), but can additionally or alternatively be performed at a remote computing system (e.g., a cloud-based system, a remotely-located server or cluster, etc.) subsequent to and/or simultaneously with (e.g., via streaming data) transmission of vehicle sensor data to the remote computing system (e.g., 520).

[0054] In some variations, at least one predetermined driving event is detected based on sensor data from any one or combination of sensors described herein, and can be performed by implementing a set of rules in the form of a model, such as an artificial neural network, as described herein. As described herein, driving event detection is preferably performed, at least in part, onboard the vehicle, but can additionally or alternatively be performed at a remote computing system subsequent to and/or simultaneously with transmission of vehicle sensor data to the remote computing system (e.g., 520).

[0055] Driving context can additionally or alternatively include the driving environment (e.g., what are the objects in the scene surrounding the vehicle, where such objects are located, properties of the objects, etc.). The driving environment can be continuously or discontinuously sensed, recorded, or otherwise suitably determined. Driving environment determination can be performed, in variations, in response to a trigger (e.g., an event-based trigger, a threshold-based trigger, a condition-based trigger etc.). In further variations, Block S100 can include iteratively recording vehicle sensor data and processing the vehicle sensor data to generate an output that can be used to trigger or otherwise suitably initiate further vehicle sensor data recordation; for example,

the method can include: continuously recording image data from an exterior-facing camera (e.g., 536) in accordance with a variation of Block S100; detecting an object in the image data in accordance with Block S200; and, recording interior and exterior image data at an interior-facing camera and the exterior-facing camera, respectively, in response to the object detection (e.g., in accordance with the variation of Block S100 and/or an alternative variation of Block S100). Collecting vehicle sensor data can include sampling at sensors of a sensor system (e.g., onboard vehicle system), receiving sensor data from the vehicle, and/or otherwise suitably collecting sensor data. Any suitable number of sensor streams (e.g., data streams) can be sampled, and sensors can be of various types (e.g., interior IMU sensors and exterior-facing cameras in conjunction, interior and exterior facing cameras in conjunction, etc.).

[0056] Block S200 includes extracting driving context data and driver behavior data from the vehicle sensor data. Block S200 functions to process the raw sensor data and derive (e.g., extract) parameters and/or characteristics that are related to the driving context and driver actions during vehicle events.

[0057] In some variations, driver behavior data includes vehicle control inputs provided by a human driver (e.g., steering, acceleration, and braking system inputs). The vehicle control inputs are preferably directly received from a vehicle control system of the vehicle, but can alternatively or additionally be inferred from the sensor data (e.g., from the external images using SLAM, from the IMU measurements, etc.). In some variations, the vehicle control inputs are directly received from an OBD (on-board diagnostic) system or an ECU (engine control unit) of the vehicle. The vehicle control inputs can be continuously obtained, or alternatively, obtained in response to detecting at least one predetermined driving event or satisfaction of a set of data sampling conditions.

[0058] In some variations, a single set of vehicle control inputs is obtained in response to detection of a driving event (e.g., steering inputs). In some variations, several sets of vehicle control inputs (e.g., steering and acceleration inputs) are obtained in response to detection of a driving event (e.g., within a predetermined time period or until a stopping condition is satisfied).

[0059] In relation to Block S200, extracting driving context data and/or driver behavior data can be performed by implementing a set of rules in the form of a model, such as an artificial neural network (e.g., a convolutional neural network), Bayesian model, a deterministic model, a stochastic and/or probabilistic model, and any other suitable model (e.g., any suitable machine learning as described above). Extracting data is preferably performed, at least in part, onboard the vehicle (e.g., at an onboard vehicle system, a vehicle computing unit, an electronic control unit, a processor of the onboard vehicle system, a mobile device onboard the vehicle, etc.), but can additionally or alternatively be performed at a remote computing system (e.g., a cloud-based system, a remotely-located server or cluster, etc.) subsequent to and/or simultaneously with (e.g., via streaming data) transmission of vehicle sensor data to the remote computing system.

[0060] Block S200 includes Block S210, which includes extracting driving context from the vehicle sensor data (e.g., sensor data provided by at least one of the sensors 531-536, shown in FIGURE 2). Block S210 functions to obtain data describing objects in the external scene (e.g., object parameters, object characteristics, object kinematics, etc.). In one variation, the driving context data can be extracted from the entire external scene captured by the external sensor data. In this variation, the driving context data can be extracted in a manner agnostic to the attention paid by the driver (e.g., irrespective of driver attention on objects as determined in one or more variations of Block S220, unweighted by region of interest / ROI, etc.), or otherwise account for driver attention. In a second variation, the driving context data can be extracted from the region of the external scene encompassed by the driver's region of interest, or be otherwise influenced by driver behavior. In a third variation, the driving context data can be extracted from a region of the external scene associated with other driving context data. For example, the driving context data can be extracted from the ahead, right, then left regions of the external scene for data sets associated with an intersection. In a second example, the driving context data can be extracted from the ahead, the front left, and the front right regions of the external scene for data sets associated with near-collision events. However, driving context data can additionally or alternatively (e.g., in a second instance of Block

S210 after generation or determination of a driving policy) be extracted based on a driving policy (e.g., taking into account a region of interest or weighting of various portions of the geospatial scene or time period of the vehicle event based on a driving policy).

[0061] In relation to Block S210, driving context data can include any data related to vehicle operation, vehicular traffic (e.g., near-miss or near-collision events; traffic operations such as merging into a lane, changing lanes, turning, obeying or disobeying traffic signals, etc.), data describing non-vehicular objects (e.g., pedestrian data such as location, pose, and/or heading; building locations and/or poses; traffic signage or signal location, meaning, pose; etc.), environmental data (e.g., describing the surroundings of the vehicle, ambient light level, ambient temperature, etc.), and any other suitable data. However, driving context data can include any other suitable data related to vehicle events, driving events, driving scenarios, and the like.

[0062] Block S210 can include performing simultaneous localization and mapping (SLAM) of the host vehicle. Mapping can include localizing the host vehicle within a three-dimensional representation of the driving context (e.g., a scene defining the positions and trajectories of the objects involved in the vehicle event).

[0063] Block S210 can include extracting object parameters from the vehicle sensor data. Object parameters can include object type (e.g., whether an object is a vehicle, a pedestrian, a roadway portion, etc.), object intrinsic characteristics (e.g., vehicle make and/or model, object shape, object size, object color, etc.)

[0064] Block S210 can include extracting vehicle event data by determining that a combination of sampled measurement values substantially matches a predetermined pattern indicative of known vehicle operational behavior (e.g., performing curve fitting on a curve of acceleration versus time curve to identify a predetermined pattern and/or a set of curve features known to correspond to a vehicle turning through a certain subtended angle). In a second variation, extracting driving context data includes translating data received from an OBD II port of the vehicle (e.g., using a lookup table). In a third variation, extracting vehicle operational data includes determining vehicle speed and direction by implementing a set of rules that track road markings and/or

landmarks in collected imagery as the markings and/or landmarks move through a sequence of image frames (e.g., using optical flow image processing, classical computer vision processing, trained machine-learning-based computer vision, etc.). In a fourth variation, extracting driving context data includes determining the location of the vehicle by combining GPS and inertial information (e.g., using IMU data used for dead-reckoning localization, using image data for extraction of inertial or motion information, etc.). In a fifth variation, extracting driving context data includes estimating a vehicle speed and/or acceleration based on microphone measurements of an audible vehicle parameter (e.g., an engine revolution parameter or revolutions per minute, a road noise parameter or decibel level of background noise, etc.). However, extracting driving context data can include otherwise suitably determining data describing agents, objects, and time-series states associated with aspects of a driving context based on collected vehicle sensor data.

[0065] Block S210 can include extracting, from exterior sensor data (e.g., image data, LIDAR data, and the like) of a driving data set, external scene features of an external scene of the vehicle (e.g., 501) represented by the exterior sensor data of a driving data set. In some variations, extracting scene features from the exterior sensor data is performed at the onboard vehicle system (e.g., 510). In some variations, the onboard vehicle system (e.g., 510) transmits the exterior sensor data to a remote computing system (e.g., 520), and the remote computing system extracts the external scene features.

[0066] In some variations, external scene features are extracted from one or more portions of the exterior sensor data that correspond to a region of interest (ROI) of the external scene of the vehicle (e.g., 501), and the features extracted from an ROI are used to train a driving response model, as described herein. Alternatively, the external scene features can be extracted from the full frame(s) of the external image(s). However, the external scene features can be extracted from any other suitable portion of the external scene and/or representation thereof.

[0067] In some variations, regions of interest of the external scene are determined at the onboard vehicle system (e.g., 510). In some variations, the onboard vehicle system (e.g., 510) transmits the exterior sensor data to a remote computing system (e.g., 520),

and the remote computing system determines regions of interest of the external scene. However, the ROI can be determined by any other suitable system.

[0068] In a first variation, one or more regions of interest of the external scene are determined based on driver attention of a human driver of the vehicle. In some variations, the driver attention is determined based on interior image data (sensed by an interior facing camera, e.g., 535) that is synchronized with the exterior sensor data. In some variations, the exterior sensor data is image data. In this manner, external scene features used to train the model can be features that correspond to features that a vehicle driver believes to be important. By filtering out external scene features based on importance to a human driver, a driving response model can be more accurately trained to emulate driving of a human.

[0069] In a second variation, one or more regions of interest of the external scene are determined based on external driving context data and/or the type of detected event (e.g., vehicle presence at an intersection, detection of a near-collision event, detection of tailgating, detection of hard braking, detection of hard steering, detection of quick acceleration, detection of a pedestrian, detection of an intended lane change, etc.). For example, in a case of the driving context data indicating the presence of the host vehicle at an intersection, the forward, right, then left regions of the external scene can be determined as regions of interest for the external scene (e.g., in sequence). As another example, in a case of driving context data indicating a forward near-collision event, an forward region of the external scene can be determined as a region of interest for the external scene. As another example, in a case of driving context data indicating a forward near-collision event, the right, then left regions of the external scene can be determined as regions of interest, thereby providing scene information that can be used to evaluate an evasive left or right turn maneuver.

[0070] In some variations, a region of interest in the external scene is identified by determining a saliency map, as described herein. In some variations, the saliency map is a collection of saliency scores associated with the external scene of the vehicle represented by the exterior sensor data (e.g., image data, LIDAR data, and the like) sampled at block

S100. In some variations, the saliency scores are associated with a driving event that corresponds to the external sensor data (e.g., a driving event that triggered collection of the external sensor data, a driving event detected at a time corresponding to a time associated with the external sensor data, etc.). In some variations, each saliency score corresponds to a spatiotemporal point with the external scene represented by the external sensor data. Variations of saliency maps are described herein with respect to block S310. In some variations, each saliency score is proportional to a duration of intensity of attention applied by the driver of the vehicle to the spatiotemporal point, as indicated by the driver's eye gaze and/or head pose (e.g., the direction of a driver's eye gaze, the angular orientation of a driver's head, etc.) extracted from the series of interior images (e.g., over time) (sensed by an interior facing camera, e.g., 535) that are synchronized with the exterior sensor data. Regions of interest (e.g., to the driver, to the model) within the external scene can be defined within the saliency map based on the saliency scores. For example, a region of the saliency map can be defined as a region of interest (e.g., associated with a particular time and/or region of space of the vehicle event) if the saliency score exceeds a threshold value at each spatiotemporal point in the region.

[0071] In some variations, each determined region of interest (ROI) of the external sensor data can be used to train a region of interest (ROI) model (e.g., attention model, scanning model) that returns a region of interest in a 3D (external) scene or a region of interest in the exterior-facing image. The ROI model can ingest a set of external scene features (e.g., wherein the external scene features can be used as the input), auxiliary sensor data, and/or any other suitable information, and output an external scene ROI, set thereof, or series thereof for analysis. The ROI model can be trained on a data set including: external scene features, auxiliary sensor data, and/or any other suitable information of driving data sets, associated with (e.g., labeled with) the ROIs determined from the respective driving data sets. In some variations, driving data sets used to train the ROI model are filtered based on an associated driving quality metric (e.g., filtered for "good" drivers or "good" driving behaviors). In some variations, a driving quality metric is determined for each candidate driving data set, wherein the candidate driving data sets

having a driving quality metric satisfying a predetermined set of conditions (e.g., higher than a threshold quality score, lower than a threshold quality score, having a predetermined quality classification or label, etc.) are selected as training data sets to be used to train the ROI model. In some variations, driving data sets used to train the ROI model are filtered on an event associated with the driving data set, and the filtered driving data sets are used to train an event-specific ROI model. However, any other suitable driving data set can be used to train the ROI model.

[0072] In some variations, block S210 is performed at the onboard vehicle system (e.g., 510); however, Block S210 can be otherwise suitably performed at any other suitable system and/or system components.

[0073] Block S200 includes Block S220, which includes extracting driver behavior data from the vehicle sensor data e.g., sensor data provided by at least one of the sensors 531-536 shown in FIGURE 2). Block S220 functions to determine actions taken by the driver while operating the vehicle (e.g., driving the vehicle, occupying the vehicle while the vehicle is running, etc.). Block S220 can also function to determine a distraction level of the driver. Block S220 can include determining driver (e.g., vehicle) control inputs (e.g., acceleration, steering, braking, signaling, etc.), determining occupant (e.g., driver and/or passenger) data, and determining a driver distraction factor (e.g., a value characterizing a level of driver distraction), each of which can be performed substantially simultaneously, asynchronously, periodically, and/or with any other suitable temporal characteristics. Block S220 is preferably performed at the onboard vehicle system (e.g., 510); however, Block S220 can be otherwise suitably performed at any other suitable system and/or system components. Block S220 is preferably performed using sensor signals (e.g., images of the vehicle interior, measurements, etc.) concurrently sampled with the signals (e.g., exterior images) from which vehicle events are extracted, but can alternatively or additionally be sampled before, after, or at any suitable time relative to the signals from which the vehicle events are extracted.

[0074] In some variations, driver behavior is determined based on vehicle behavior (e.g., hard braking, hard steering, fast acceleration, erratic driving behavior, etc.). In a

first example, vehicle control inputs of a driver can be inferred without receiving explicit vehicle control inputs provided by the driver. Instead, vehicle behavior (such as movement of the vehicle, activation of vehicle turn signals, etc.) as determined by vehicle sensor data, can be used to infer control of the vehicle by the driver. For example, if sensor data indicates that the vehicle is moving left, a steer-left vehicle control input provide by the driver can be inferred. In another example, movement of the vehicle can be characterized (e.g., sudden stopping, sudden turning, fast acceleration, erratic movement, etc.) based on vehicle sensor data, and the movement of the vehicle can be used to determine a driving behavior of the driver (e.g., hard braking, hard steering, fast acceleration, erratic driving behavior, etc.).

[0075] Block S220 can include can include extracting interior activity data. Extracting interior activity data includes extracting data from a data stream (e.g., an image stream, a gyroscopic data stream, an IMU data stream, etc.) that encodes information concerning activities occurring within a vehicle interior. Such interior activity data can include driver activity (e.g., driver gaze motion, driver hand positions, driver control inputs, etc.), passenger activity (e.g., passenger conversation content, passenger speaking volume, passenger speaking time points, etc.), vehicle interior qualities (e.g., overall noise level, ambient light level within the cabin, etc.), intrinsic vehicle information perceptible from within the vehicle (e.g., vibration, acoustic signatures, interior appointments such as upholstery colors or materials, etc.), and any other suitable activity data related to the vehicle interior and/or collected from within the vehicle (e.g., at the onboard vehicle system).

[0076] In variations, determining driver behavior can include determining (e.g., via gaze direction analysis of the vehicle sensor data) the driver's comprehension of the vehicle surroundings during the vehicle event and correlating the driver's attention to various portions of the surroundings with the dynamics of the vehicle event (e.g., via saliency mapping).

[0077] In variants, determining the driver gaze direction can be difficult because there is no ground truth to determine which object or scene region of interest that the

user is gazing at (e.g., because this would require drivers to label the scene region of interest while driving). In one embodiment, the method can extract the driver's eye gaze and/or head pose (e.g., the direction of a driver's eye gaze, the angular orientation of a driver's head, etc.) from a series of interior images (e.g., over time) to infer scanning patterns from the driver. The scanning pattern can be used to determine the range of the scene (e.g., from the extremities of the scanning pattern), or otherwise used (e.g., to determine whether the driver is scanning a region of interest that a region of interest model, trained on a simulation or historic data from one or more drivers, identifies). The system can further use the scanning pattern to infer the regions (regions of interest) in the external scene that the user is looking at, based on the gaze direction (determined from the interior-facing camera) and the camera calibration parameters (e.g., relating the interior and exterior cameras, such as the extrinsic matrix). These regions can optionally be used to generate a training dataset, which can include: interior images annotated for gaze, and exterior images annotated for region of interest (e.g., the region that the driver was looking at).

[0078] Driver and/or operator behavioral data can include: operator profiles (e.g., history, driver score, etc.); operator behavior (e.g., user behavior), such as distraction level, expressions (e.g., surprise, anger, etc.), responses or actions (e.g., evasive maneuvers, swerving, hard braking, screaming, etc.), cognitive ability (e.g., consciousness), driving proficiency, willful behavior (e.g., determined from vehicle control input positions), attentiveness, gaze frequency or duration in a predetermined direction (e.g., forward direction), performance of secondary tasks (e.g., tasks unrelated to driving, such as talking on a cell phone or talking to a passenger, eating, etc.), or other behavior parameters; or any other suitable operator parameter.

[0079] Block S220 can include determining an intended action of a driver. For example, Block S220 can include determining that the driver intends to change lanes based on the driver performing a series of actions including scanning a region to the side of the vehicle (e.g., in the lane-change direction), checking a blind spot to the side of the vehicle (e.g., in the lane-change direction), and other suitable preparatory actions related

to lane changing. In another example, Block S220 can include determining that a driver intends to brake based on a decision tree populated with possible actions based on vehicle event data (e.g., extracted in accordance with one or more variations of Block S210). In a third example, the intended actions can be determined from navigational systems (e.g., a driving directions client or application).

[0080] Block S220 can optionally include determining one or more driving actions of the driver. The driving actions can be associated with the vehicle event, the driver behaviors, and/or any other suitable information, wherein the associated dataset can be used to train the driving policy model(s) and/or any other suitable models. The driving actions are preferably the actions that the driver takes in response to the vehicle event, but can alternatively be actions that the driver takes before and/or after the vehicle event. In one example, the driving actions can be a driving maneuver (e.g., right turn, left turn, reverse, driving straight, swerving, etc.) that the driver took after the vehicle event (e.g., arrival at an intersection, occurrence of a near-collision event, etc.). However, the driving actions can be otherwise defined. The driving actions can be determined from: the interior images, the exterior images, vehicle controls (e.g., determined from the CAN bus, etc.), vehicle ego-motion, signals sampled by an exterior sensor (e.g., a sensor on a second vehicle), or otherwise determined.

[0081] Block S220 can include determining an attention level of a driver associated with an object described by the vehicle event data. For example, Block S220 can include calculating the time duration that a driver has directed his or her gaze at an object or region present in the vehicle surroundings (e.g., a cross-street outlet, a secondary vehicle, a traffic signal, etc.). However, Block S220 can include otherwise suitably determining a driver's attention level in any other suitable manner. Determining the attention level can function to provide an input for determining a saliency score for a point or region in space during a driving event.

[0082] In some variations, block S200 includes determining a driving quality metric for a driving data set. In some variations, determining a driving quality metric is performed at the onboard vehicle system (e.g., 510). In some variations, the onboard

vehicle system (e.g., 510) transmits a driving data set to a remote computing system (e.g., 520), and the remote computing system determines a driving quality metric for the driving data set.

[0083] In some variations, a driving quality metric includes one or more of a driver attentiveness score, a maneuver skill metric, an overall skill metric for a driving session of a driving data set, an overall skill (e.g., across multiple driving sessions) of a driver as identified in a profile, and a ride comfort metric.

[0084] In some variations, at least one of the onboard vehicle system 510 and the remote system 520 determines a driver attentiveness score. In a first variation, a region of interest model is used to determine the attentiveness score. The attentiveness score is determined by determining a region in a scene represented by the exterior image data that the driver (of the driving data set) is looking at based an eye gaze of the driver extracted from the interior image data, uses a region of interest (ROI) model (as described herein) to identify at least one region of interest from the exterior sensor (e.g., image) data, and determines the driver attentiveness score by comparing a region the driver is looking at to regions determined by the ROI model.

[0085] In a second variation, the driver attentiveness score is determined by using commonly accepted driving standards. The driver attentiveness score is determined by determining a region in a scene represented by the exterior image data that the driver (of the driving data set) is looking at based an eye gaze of the driver extracted from the interior image data, and comparing where the driver is looking to what the driver should be looking at according to commonly accepted driving standards. In one example, the driver attentiveness can be determined using the methods disclosed in US Application 16/239,326 filed 03-JAN-2019, incorporated herein in its entirety by this reference.

[0086] In some variations, determining a region in a scene represented by the exterior image data that the driver (of the driving data set) is looking at determining (e.g., via gaze direction analysis of the interior image data) the driver's comprehension of the vehicle surroundings and correlating the driver's attention to various portions of the

scene represented by the exterior sensor data with the dynamics of the vehicle event (e.g., via saliency mapping).

[0087] Determining driver attention can include calculating the time duration that a driver has directed his or her gaze at an object or region present in the exterior scene (e.g., a cross-street outlet, a secondary vehicle, a traffic signal, etc.). Determining the attention level can function to provide an input for determining a saliency score for a point or region in space (represented by the exterior sensor data) during a driving event.

[0088] In some variations, the driver's eye gaze and/or head pose (e.g., the direction of a driver's eye gaze, the angular orientation of a driver's head, etc.) is extracted from a series of interior images (e.g., over time) (sensed by an interior facing camera, e.g., 535) that are synchronized with the exterior sensor data to infer scanning patterns from the driver. The range of the external scene can be determined from the extremities of the scanning pattern. The scanning pattern can be used to infer the regions (regions of interest) in the external scene that the user is looking at, based on the gaze direction (determined from the interior-facing camera, e.g., 535) and the camera calibration parameters (e.g., relating the interior camera 535 and exterior camera 536, such as the extrinsic matrix). These regions can optionally be used to generate a driving dataset for a vehicle (as described herein), which can include: interior images annotated for gaze, and exterior images annotated for region of interest (e.g., the region that the driver was looking at).

[0089] In some variations, determining a region in a scene represented by the exterior image data that the driver (of the driving data set) is looking at based on an eye gaze of the driver extracted from the interior image data includes: determining whether the driver is looking at locations of high saliency based on detecting the driver's gaze in the interior image data.

[0090] The attentiveness score can be determined in response to detection of a vehicle event, and the attentiveness score can indicate whether the human driver is looking in a direction of high saliency for the detected vehicle event. For example, if a driver is intending to make a right turn at an intersection, while the exterior-facing

camera captures the other object trajectories in relation to the path of the host vehicle, the interior-facing camera can capture whether the driver was looking in the correct directions to suitably execute the right turn. A trained region of interest (ROI) model (as described herein) can be used to determine the locations of high saliency in external sensor data (e.g., image data, point cloud) that represents an external scene. Determining the attentiveness score can be performed in real-time during vehicle operation, or subsequent to a driving action performed by a human driver.

[0091] At least one of the onboard vehicle system 510 and the remote system 520 determines a skill of a maneuver associated with the driving data set. In a first variation, a trained driving policy model is used to determine the skill of a maneuver. The skill of a maneuver is determined by comparing a driver's driving actions of the maneuver (identified for the driving data set) with driving actions determined by a trained driving policy model (as described herein) from the synchronized interior sensor data and exterior sensor data of the driving data set.

[0092] Determining a skill of a maneuver by using a trained driving policy model can optionally include: the driving response model receiving a region of interest (ROI) for a scene represented by exterior sensor data and the external sensor data; and outputting a driving action for the maneuver to be compared with the driver's driving actions of the maneuver (identified for the driving data set). In some variations, the trained ROI model receives the exterior sensor data and the interior image data as inputs, and the outputs the ROI for the scene based on these inputs.

[0093] In a second variation, a skill of a maneuver associated with the driving data set is determined by comparing a driver's driving actions of the maneuver (identified for the driving data set) with commonly accepted driving standards.

[0094] In some variations, a driver skill of the driver of the driving data is determined based on a stored driver profile that specifies a skill of the driver. The driver profile can be stored at the remote computing system 520. A driver ID of the driver is used to retrieve the stored driver profile and identify the drive skill. In a first variation, at least one of the onboard vehicle system 510 and the remote system 520 determines a

driver ID of the driver of the driving data set based on the interior image data. At least one of the onboard vehicle system 510 and the remote system 520 can tag the driving data set with the determined driver ID. At least one of the onboard vehicle system 510 and the remote system 520 can tag the driving data set with the retrieved driver skill.

[0095] In a third variation, the driving quality score can be determined based on a ride comfort metric of a driving session associated with a driving data set. The ride comfort metric can be determined based on sensor data included in the driving data set (e.g., lateral acceleration, G forces, etc.), or otherwise determined. In this manner, a driving policy model can be trained to generate driving actions that result in an autonomous vehicle driving in a comfortable manner, as opposed to a jerky manner that might cause discomfort to a passenger.

[0096] Block S300 includes determining a driving response model (e.g., driving policy model). S300 functions to generate a model that outputs naturalistic driving behaviors (e.g., driving policy). S300 is preferably performed by the remote computing system, but can alternatively be performed at the onboard vehicle system (e.g., 510) or at any other suitable system. In some variations, the onboard vehicle system (e.g., 510) transmits a driving data set (and, optionally, external scene features extracted from the driving data set) to a remote computing system (e.g., 520), and the remote computing system determines a driving policy by using the driving data set.

[0097] In some variations, determining a driving policy includes training a driving response model based on the extracted vehicle event data and driver behavior data.

[0098] The driving response model is preferably a convolutional neural network, but can alternatively be a deep neural network, a Bayesian model, a deterministic model, a stochastic and/or probabilistic model, a rule-based model, and any other suitable model or combination thereof. The driving response model is preferably trained using reinforcement learning, but can alternatively be trained using supervised learning, unsupervised learning, or otherwise trained.

[0099] In some variations, block S300 includes accessing driving data sets for a plurality of human-driven vehicles. Each driving data set includes synchronized interior

sensor data and exterior sensor data, and vehicle control inputs associated with the synchronized interior sensor data. The interior sensor data can include interior image data. The exterior sensor data can include 2-D image data, 3-D image data, a point cloud (e.g., LIDAR output data), and the like. The driving data sets can include driving data sets for a plurality of vehicles, driving sessions, drivers, driving contexts, and the like. In some variations, one or more driving data sets include information indicating at least one of the following for the driving data set: a driving event, a vehicle identifier, a driver identifier, a driving session, driving context information, a driving quality metric, a driver score, a driver skill, a driver profile, and the like.

[00100] In some variations, block S300 includes selecting (from a plurality of driving data sets) driving data sets having a driving quality metric that satisfies a predetermined set of conditions, and training a driving response model based on external scene features (block S210) extracted from the selected driving data sets and the vehicle control inputs from the selected driving sets.

[00101] In some variations, the predetermined conditions can include driver performance score thresholds indicating a level of driver skill (e.g., an overall driver skill, a driver skill for a driving session, a skill for a particular maneuver of the driving data set, and the like). A driver skill can be assigned to each maneuver in a driving data set, and the driving data set can be selected for use in training based on the skills assigned to the maneuvers, or alternatively, portions of a driving data set corresponding to maneuvers having skill levels above a threshold can be selected for use in training. A driving data set for a driving session for a driver having an overall driving skill (as defined in a driver profile) can be selected for use in training. A driving data set for a driving session for a driver having a driving skill (determined for the driving session of the driving data set) can be selected for use in training, regardless of an overall driver skill for the driver as determined from previous driving sessions.

[00102] In some variations, predetermined conditions can include attentiveness score thresholds indicating a level of driver attentiveness.

[00103] In some variations, predetermined conditions can include ride comfort thresholds.

[00104] In some variations, predetermined conditions can include driving context features. For example, the method can include identifying data sets associated with a predetermined driving event; identifying data sets associated with a predetermined ambient driving environment; identifying data sets associated with a predetermined time of day; identifying data sets associated with a predetermined location or location class (e.g., intersection, freeway, traffic, etc.); or any other suitable set of driving context features.

[00105] Block S300 includes selecting driving data sets based on driving quality metric. In some variations, block S300 is performed at the onboard vehicle system (e.g., 510). In some variations, the onboard vehicle system (e.g., 510) transmits a driving data set to a remote computing system (e.g., 520), and the remote computing system selects the driving data sets based on driving quality metric.

[00106] In some variations, driving data sets used to train the driving response model are filtered on an event associated with the driving data set, and the filtered driving data sets are used to train an event-specific driving response model.

[00107] In some variations, block S300 functions to determine a driving policy based on the driving context data in combination with the driver behavior data. In some variations, Block S300 functions to convert driver behavior in various driving scenarios (e.g., determined in accordance with one or more variations of Block S200) to a set of context-based decision-making rules that collectively define a driving policy. The driving policy can be implemented as a model (e.g., a set of explicitly programmed rules that process the inputs of vehicle event data and output a set of desirable driver behavior, a trained or trainable machine-learning model, a combination of probabilistic and deterministic rules and/or parametric equations, etc.) or otherwise suitably implemented.

[00108] In relation to Block S300, the driving policy can be generated based on the driver actions during a vehicle event. For example, a vehicle event can include making a

turn at an intersection, and the driver actions can include the locations and/or regions towards which the driver is looking over the time period leading up to and including making the turn. The exterior-facing camera of the onboard vehicle system can simultaneously capture the trajectories of secondary vehicles traversing the roadway proximal the intersection, and Block S300 can include first, determining that the right turn was skillfully executed (e.g., efficiently and safely executed) and second, designating the regions that received the driver's attention during the vehicle event as regions of interest and also designating the driver control inputs in relation to the driving maneuver as components of the driving policy.

[00109] Block S300 can include Block S310, which includes: determining a saliency map of the vehicle event based on the driver behavior data. Block S310 functions to associate a saliency (e.g., saliency score, saliency metric, relative saliency, absolute saliency, etc.) with each spatiotemporal component of a vehicle event, for use in determining a driving policy for the vehicle event.

[00110] Block S310 can include determining a saliency score corresponding to a spatiotemporal point, and defining a saliency map as the collection of saliency scores associated with the vehicle event. The spatiotemporal point can be a physical region of finite or infinitesimal extent, and can be defined over any suitable time duration (e.g., an instantaneous point in time, a time period within a vehicle event, a time period including the entire vehicle event, etc.). The saliency score (e.g., saliency metric) can be a relative score (e.g., normalized to any suitable value, such as 1.0, representing the peak saliency score defined in the saliency map of the vehicle event), an absolute score (e.g., defined proportional to the duration and intensity of attention applied by the driver to the spatiotemporal point during the vehicle event), and/or otherwise suitably defined.

[00111] The saliency map can be: an array of saliency metric values (e.g., for each sub-region identifier), a heat map (e.g., stored or visualized as a heat map, as shown in FIGURE 3), an equation, or be otherwise structured. The saliency map(s) or parameters thereof (e.g., factor values, weights, geolocations, etc.) can be stored temporarily (e.g., long enough to analyze the instantaneous saliency of a driver action or behavior), for the

vehicle event duration, for the driving session duration, for longer than the driving session, or for any suitable time period. All or a subset of the generated saliency maps or parameters thereof can be stored. The saliency maps (or parameters thereof) can be stored in association with a vehicle identifier or characteristic (e.g., a map of the visibility of the vehicle surroundings by a driver situated within the vehicle), geographic location or region identifier, operator identifier, vehicle kinematics, or any other suitable factor values. The saliency map can be a physical mapping (e.g., to points in physical space) and/or a conceptual mapping (e.g., to objects identified and tracked in the context of the vehicle event), or otherwise suitably defined.

[00112] As shown in FIGURE 4, Block S310 can include defining a region of interest within the saliency map based on the saliency score(s). For example, a region of the saliency map of the vehicle event can be defined as a region of interest (e.g., associated with a particular time and/or region of space of the vehicle event) if the saliency score exceeds a threshold value at each spatiotemporal point in the region. In examples, the driving policy can generate the region(s) of interest as output(s) based on the context of the vehicle (e.g., defined in the vehicle event data).

[00113] Block S300 can include Block S320, which includes training a driving policy model based on the vehicle event data in combination with the driver behavior data. Block S320 functions to codify the driver behavior in various driving contexts (e.g., vehicle events), across one or more drivers, into a driving policy model that can be applied to similar driving contexts (e.g., vehicle events) occurring elsewhere and/or at a different time.

[00114] In examples, the synchronized interior-exterior data (e.g., collected in accordance with one or more variations of Block S100) can be used to train models (e.g., subsequent to suitable factorization into inputs in accordance with one or more variations of Block S200) via a training or learning process to generate vehicle control models that function to maneuver the vehicle in the presence of multiple objects in complex driving scenarios (e.g., such as intersections). For example, the method can include determining whether a driver (e.g., a skilled driver, highly-scored driver) was looking in the correct

direction(s) during a particular driving maneuver, and training a model to heavily weight the regions corresponding to the directions in which the driver was looking in similar driving maneuvers in similar contexts. In a second example (specific example shown in FIGURE 6A), the method can include: identifying data associated with both vehicle events of interest (e.g., data associated with intersection geotags, near-collision labels, etc.), wherein the data can include the vehicle event data (e.g., driving context data, such as external objects, object poses, vehicle kinematics, intended next action, etc.), the associated (e.g., contemporaneous) driver behavior (e.g., gaze direction, etc.), and the driver action (e.g., actions taken at the intersection, in response to the near-collision event, etc.); extracting features from the dataset (e.g., vehicle event features, driver behavior features, driver action features); and training one or more models based on the extracted features. The data can optionally be filtered for data generated by “good” drivers (e.g., driving data generated by drivers with a driver score above a predetermined threshold, drivers having a skill level or score above a predetermined threshold associated with a driver profile, etc.), but can optionally include any other suitable driving data.

[00115] In a specific example, the vehicle event features and the driver behavior features can be used to train a region of interest model that returns a region of interest in a 3D (external) scene or a region of interest in the exterior-facing image, given a set of vehicle event features (e.g., wherein the vehicle event features can be used as the input, and the driver behavior features can be used as the desired output in the labeled training set).

[00116] In a second specific example, the region(s) of interest (regions of the exterior images, regions of the 3D external scene, etc.), optionally obstacle features of the obstacles detected in the region(s) of interest, and the driver actions can be used to train a driving policy model that returns a driving action, given a set of regions of interest and/or obstacle features (specific example shown in FIGURE 6B). The driving action can be used to: control vehicle operation, score human drivers, or otherwise used. The region(s) of interest can be a single frame or a series of regions. Image processing, object detection, pose detection, trajectory determination, mapping, feature extraction, and/or

any other suitable data can be determined using conventional methods, proprietary methods, or otherwise determined. The model thus trained can then control a host vehicle to execute driving maneuvers by analyzing the vehicles that were of interest (e.g., based on the heavily weighted regions, which are chosen based upon what the human driver was monitoring in similar maneuvers).

[00117] Block S320 can include generating a training dataset based on the driving policy (e.g., generated using a driving policy model). The training dataset can include synchronized interior-exterior data annotated based on the driving policy (e.g., wherein the exterior imagery is annotated with a saliency map or saliency scores, and the interior imagery is annotated with gaze direction), and/or any other suitable data for training a computational model based on the driving policy.

[00118] In a variation, Block S320 includes training a driving policy model embodied as a computational learning network (e.g., a convolutional neural network) for vehicle control using vehicle event data weighted by driver attention (e.g., according to a saliency map of each vehicle event determined in accordance with one or more variations of Block S310), such that the network learns to focus on certain regions of interest in the vehicle event data. The regions of interest can be physical regions (e.g., geolocation regions) within a scene depicted in vehicle sensor data, temporal regions of interest (e.g., time periods of interest during vehicle events), and/or otherwise suitably defined.

[00119] The method can include Block S400, which includes controlling a vehicle based on the driving policy (example shown in FIGURE 6B). Block S400 functions to implement the driving policy in conjunction with vehicle operation (e.g., autonomous operation, semi-autonomous driving, etc.). Controlling the vehicle based on the driving policy can include implementing the decision output of a driving policy model (e.g., developed based on decisions made by a human driver, by an autonomous control system, etc.) in a driving scenario where there are multiple objects with associated dynamics and behaviors provided to the driving policy model as inputs.

[00120] Block S400 can include examining the vehicle surroundings (e.g., objects in the scene imaged by the onboard vehicle system) to determine a preferred option for

navigating the vehicle (e.g., among candidate navigation routes) amidst the agents (e.g., obstacles, objects, etc.) occupying the vehicle surroundings, and controlling the vehicle in real-time based on the determined navigation option (e.g., route).

[00121] In some variations, controlling an autonomous vehicle includes: determining external scene features from a set of external scene information S410, providing the external scene features to a driving response model (e.g., trained using the methods described above) to determine vehicle control inputs for the autonomous vehicle S420, and controlling the autonomous vehicle based on the vehicle control inputs S430. In some variations, the driving response model is trained as described herein with respect to block S300. In some variations, the driving response model is trained on historic driving data sets for historic human-driven vehicles. In some embodiments, the historic driving data sets are constructed as described herein in relation to blocks S100 and S200. In some variations, the historic driving data sets are associated with driving quality metrics satisfying predetermined set of conditions, as described herein in relation to block S300. In some variations, historic driving data sets include historic external scene features extracted from historic external images, and historic vehicle control inputs associated with the historic external images.

[00122] In some variations, the driving response model is specific to a driving event, and auxiliary sensor data (as described herein) of the vehicle is monitored for occurrence of the driving event. In response to occurrence of the driving event, the driving response model is selectively executed, the external scene features are selectively provided to the driving response model, and the outputs of the driving response model (e.g., control inputs) can be fed into the autonomous vehicle control model and/or used to control autonomous vehicle operation.

[00123] In some variations, auxiliary sensor data is provided to a scanning model (region of interest model) that determines a region of interest (ROI) in an external scene represented by the external scene information, and the external scene features are extracted from the determined ROI. The external scene features (and/or features from other external scene regions) are then fed to the driving response model. In a specific

example, when features from both the ROI and the other regions are fed to the driving response model, the features from the ROI can be higher-resolution, more accurate, higher-weighted, preferentially analysed, or otherwise differ from the other features. Alternatively, the features from the ROI can be treated equally as features from the other regions, or otherwise treated. In some variations, the scanning model is trained on historic regions of interest in external scenes corresponding to historic driver gaze directions, historic auxiliary sensor data associated with interior images, and historic external scene features of the external scenes. In some variations, the historic driver gaze directions are each extracted from an interior image. In some variations, the historic external scene features are features that have been extracted from external images contemporaneously sampled with the respective interior images.

[00124] In some variations, the external scene information includes external scene measurements. In some variations, the external scene measurements include LIDAR measurements. In some variations, the external scene information includes external camera image data.

[00125] The method can include Block S500, which includes providing output to a human driver based on the driving policy. Block S500 functions to apply a driving policy to driver behavior and suggest driver actions to a human driver according to the driving policy.

[00126] Block S500 can include coaching a human driver based on the driving policy. For example, Block S500 can include providing an output that coaches the driver to scan the road, to look towards a region of interest (e.g., as output by a driving policy model based on a current vehicle event and/or driving scenario), to check the vehicle mirrors (e.g., side view mirrors, rear view mirrors, etc.), to be alert for pedestrians (e.g., at a crosswalk, at a surface street intersection, etc.). The output can be provided in the form of an audio output (e.g., a voice message, a beep pattern, a coded audio signal, etc.), a visual output (e.g., a rendering of a region of interest on a heads up display, an arrow pointing towards a designated region of interest rendered at a display within the vehicle, a light emitter of the onboard vehicle system, etc.), and any other suitable output type.

[00127] In variations, Block S500 can include capturing, at an interior-facing camera, whether a driver is looking at locations of high saliency (e.g., corresponding to a high saliency score, a saliency metric above a threshold, etc.) associated with a given vehicle event. For example, if a driver is intending to make a right turn at an intersection, while the exterior-facing camera captures the other object trajectories in relation to the path of the host vehicle (e.g., vehicle trajectories, pedestrian trajectories, etc.), the interior-facing camera can capture whether the driver was looking in the correct directions (e.g., towards regions of high saliency, towards directions determined based on the driving policy, etc.) to suitably execute the right turn.

[00128] Block S500 can be performed in real-time (e.g., near real-time, substantially real-time, etc.) during vehicle operation. For example, Block S500 can include alerting a human driver that the human driver is checking their blind spot at an inadequate frequency, according to the determined driving policy, based on real-time extraction of driver behavior including the human driver's gaze direction (e.g., whether the gaze direction is aligned with a rear-view and/or side-view mirror at a predetermined frequency, in relation to detected vehicle maneuvers, etc.). Additionally or alternatively, Block S500 can be performed subsequent to a driving action performed by a human driver (e.g., after the conclusion of a vehicle event, after the conclusion of a driving session, etc.). For example, Block S500 can include determining a performance score associated with a driving session and/or driver actions during a specific vehicle event (e.g., based on a comparison of driver behavior with the driving policy), and providing the performance score to the human driver subsequent to the driving session and/or vehicle event (e.g., as at least a part of a summary report of driver performance). However, Block S500 can additionally or alternatively be performed with any suitable temporal characteristics (e.g., prior to a driving session as a reminder of past performance, periodically during a driving session, periodically at any suitable frequency, continuously, asynchronously, etc.).

[00129] The method of the preferred embodiment and variations thereof can be embodied and/or implemented at least in part as a machine configured to receive a computer-readable medium storing computer-readable instructions. The instructions are

preferably executed by computer-executable components preferably integrated with a suitable system and one or more portions of a processor or controller. The computer-readable medium can be stored on any suitable computer-readable media such as RAMs, ROMs, flash memory, EEPROMs, optical devices (CD or DVD), hard drives, floppy drives, or any suitable device. The computer-executable component is preferably a general or application specific processor, but any suitable dedicated hardware or hardware/firmware combination device can alternatively or additionally execute the instructions.

[00130] Although omitted for conciseness, the preferred embodiments include every combination and permutation of the various Blocks of the method, any of which can be utilized in any suitable order, omitted, replicated, or otherwise suitably performed in relation to one another.

[00131] As a person skilled in the art will recognize from the previous detailed description and from the figures and claims, modifications and changes can be made to the preferred embodiments of the invention without departing from the scope of this invention defined in the following claims.

CLAIMS

What is claimed is:

1. A method, comprising:
 - determining a driving data set for each of plurality of human-driven vehicles, each driving data set comprising:
 - synchronized interior image data and exterior image data; and
 - vehicle control inputs associated with the synchronized interior image data and exterior image data;
 - for each driving data set, extracting, from the exterior image data, exterior scene features of an exterior scene of the respective vehicle;
 - for each driving data set, determining a driving quality metric;
 - selecting driving data sets, having a driving quality metric satisfying a predetermined set of conditions, from the plurality of driving data sets; and
 - training a driving response model based on the exterior scene features and the vehicle control inputs from the selected driving data sets.
2. The method of Claim 1, further comprising:
 - for each driving data set:
 - determining driver attention based on the interior image data; and
 - determining a region of interest (ROI) within the exterior scene based on driver attention; and
 - training an ROI scanning model based on the exterior scene features and the ROI from the selected driving data sets.
3. The method of Claim 2, wherein the driving response model is trained based on the exterior scene features extracted from the ROI.
4. The method of Claim 1, wherein a driving quality metric for a driving data set is determined based on the interior image data of the driving data set.

5. The method of Claim 4, wherein a driving quality metric of a driving data set comprises an attentiveness score, determined based on driver gaze, wherein the driver gaze is extracted from the interior image data of the driving data set.
6. The method of Claim 5, wherein the attentiveness score is determined based on the driver gaze and a region of interest identified in the exterior scene by using a scanning model that determines a region of interest in the exterior scene from the exterior image data.
7. The method of Claim 1, wherein a driving quality metric for a driving data set comprises a driver score for the respective human.
8. The method of Claim 7, wherein a driving quality metric for a driving data set is determined based on a comparison of driver behavior, indicated by the driving data set, with an expected driving behavior determined by providing at least the exterior image data of the driving data set to a trained driving response model.
9. The method of Claim 1, wherein a driving quality metric of a driving data set comprises a ride comfort metric.
10. The method of Claim 1, wherein determining the driving data set comprises, at each of the plurality of human-driven vehicles:
 - detecting a predetermined driving event; and
 - recording the synchronized interior image data and exterior image data after detecting the predetermined driving event.
11. The method of Claim 10, further comprising, at each of the plurality of human-driven vehicles, transmitting the synchronized interior image data and exterior image data to a remote computing system, wherein the remote computing system extracts the exterior scene features, determines the driving quality metric, selects the driving data sets, and trains the driving response model.
12. The method of Claim 11, wherein training the driving response model comprises:
 - selecting driving data sets for each of a plurality of different driving events; and
 - training a different driving response model for each of the plurality of different driving events by using the corresponding selected driving data sets.

13. The method of Claim 1, wherein extracting exterior scene features from the exterior image data comprises:
 - determining a point cloud based on the exterior image data; and
 - extracting the exterior scene features from the point cloud.
14. A method for autonomous vehicle control, comprising, at the autonomous vehicle:
 - determining exterior scene features from a set of exterior scene measurements;
 - determining vehicle control inputs by providing the exterior scene features to a driving response model, wherein the driving response model is trained on:
 - historic driving data sets for historic human-driven vehicles, the historic driving data sets associated with driving quality metrics satisfying a predetermined set of conditions, wherein the historic driving data sets comprise:
 - historic exterior scene features extracted from historic exterior images; and
 - historic vehicle control inputs associated with the historic exterior images; and
 - controlling the autonomous vehicle based on the vehicle control inputs.
15. The method of Claim 14, wherein the driving response model is specific to a driving event, the method further comprising:
 - monitoring sensor data of the vehicle for occurrence of the driving event; and
 - in response to occurrence of the driving event, selectively providing the exterior scene features to the driving response model.
16. The method of Claim 14, further comprising:

- providing auxiliary sensor data to a scanning model that determines a region of interest (ROI) in the exterior scene, wherein the scanning model is trained on:
 - historic regions of interest in exterior scenes corresponding to driver gaze directions, wherein the driver gaze directions are each extracted from an interior image associated with the respective exterior scene;
 - historic auxiliary sensor data associated with the respective interior image; and
 - historic exterior scene features of the exterior scenes, extracted from exterior images contemporaneously sampled with the respective interior images;
 - wherein the exterior scene features are extracted from the determined ROI.
17. The method of Claim 14, wherein the exterior scene measurements comprise LIDAR measurements.
18. A system, comprising:
- an interior camera configured to generate interior image data of a vehicle;
 - an exterior camera configured to generate exterior image data of the vehicle; and
 - a processing system configured to:
 - generate a driving data set for a driving session that includes synchronized interior image data and exterior image data provided by the interior camera and the exterior camera, respectively, and vehicle control inputs associated with the synchronized interior image data and exterior image data;
 - determine a driving quality metric for the driving data set;
 - select the driving data set based on the driving quality metric; and
 - train a driving response model based on: external scene features extracted from the exterior image data and driver behaviors extracted from the interior image data from the selected driving data set.

19. The system of claim 18, further comprising:
 - a scanning model configured to determine a region of interest (ROI) in an exterior scene from exterior image data of the exterior camera; and
 - a gaze detector configured to extract an eye gaze of a driver from the interior image data of the interior camera,wherein the onboard vehicle system is configured to use the ROI and the eye gaze to determine the driving quality metric for the driving data set.
20. The system of claim 18, further comprising:
 - a trained driving response model configured to determine a driving behavior from at least the exterior image data of the exterior camera,wherein the onboard vehicle system is configured to generate a notification when the driving behavior differs from the vehicle control inputs.

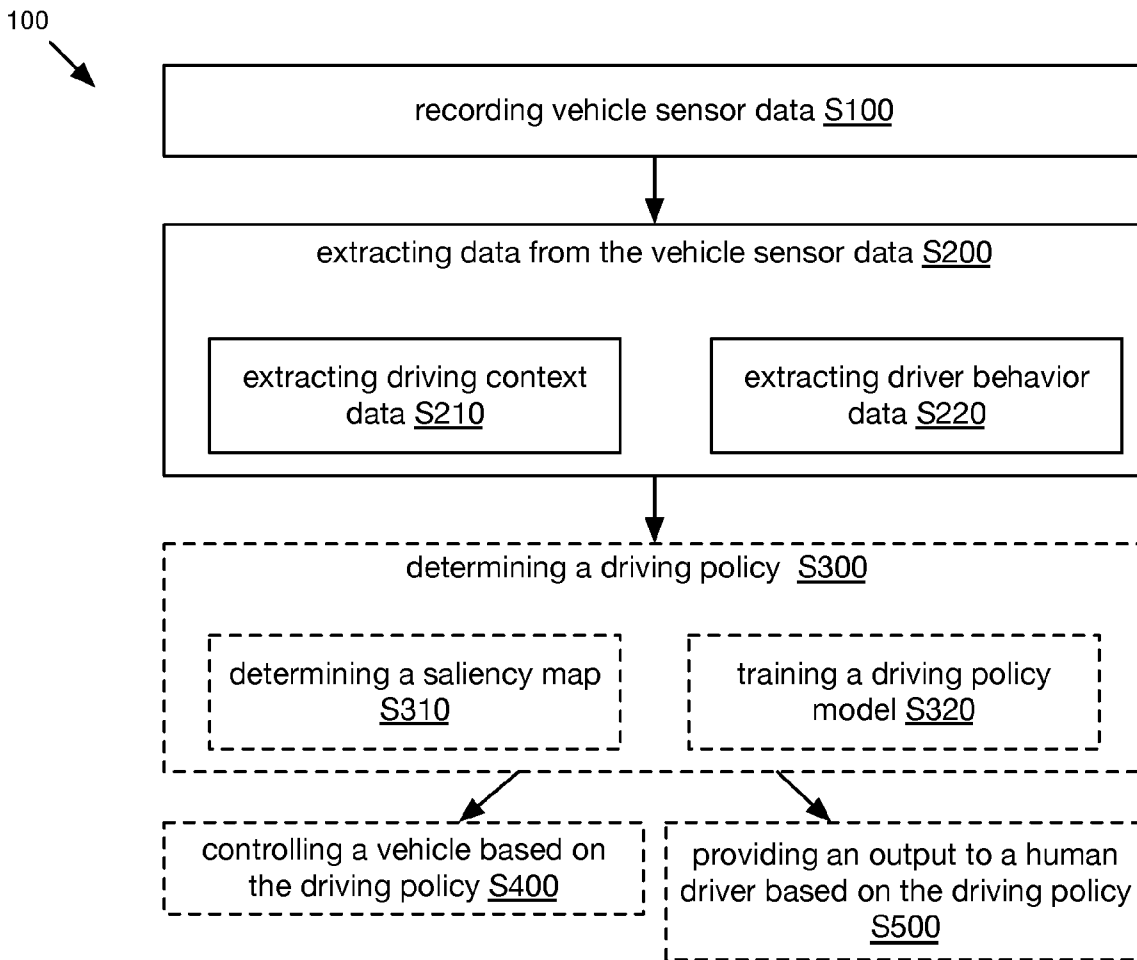


FIGURE 1




FIGURE 2

3 / 8




FIGURE 3

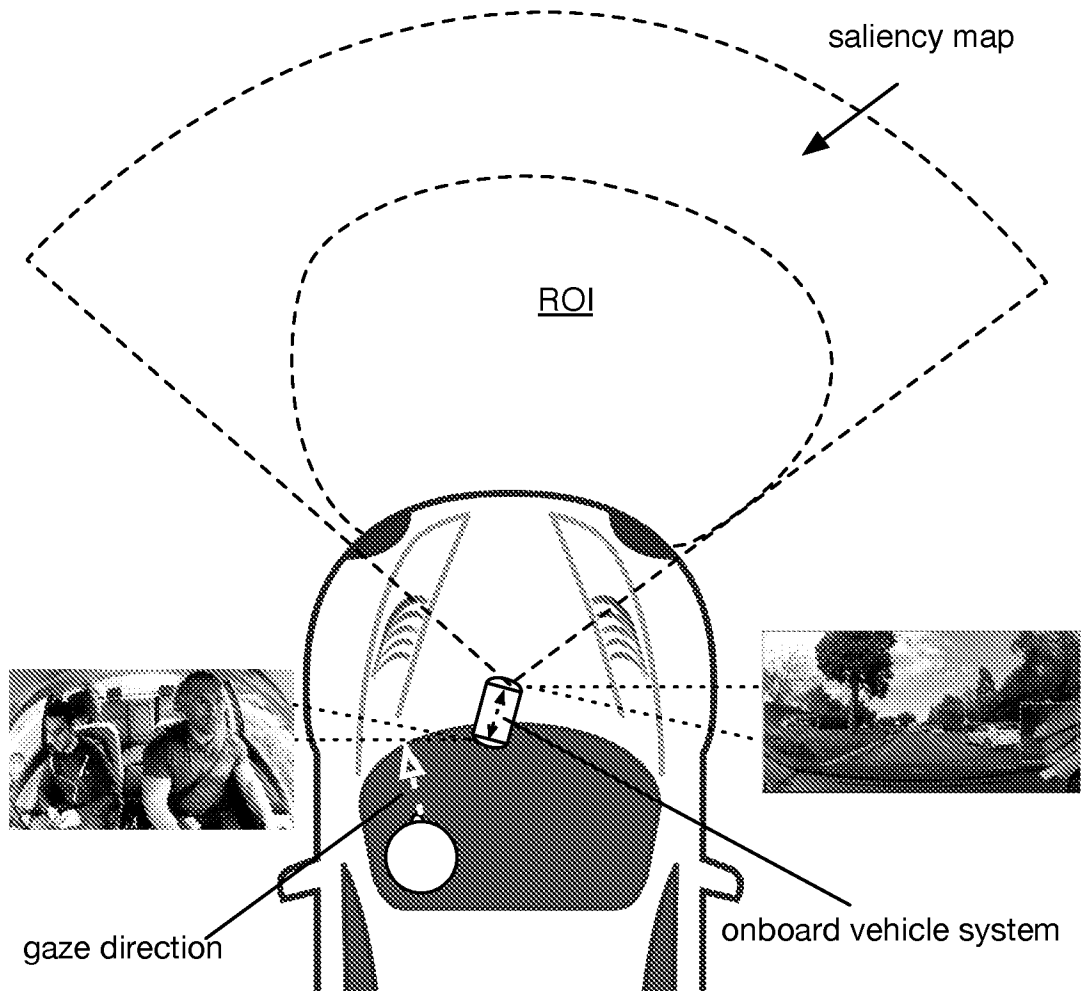


FIGURE 4

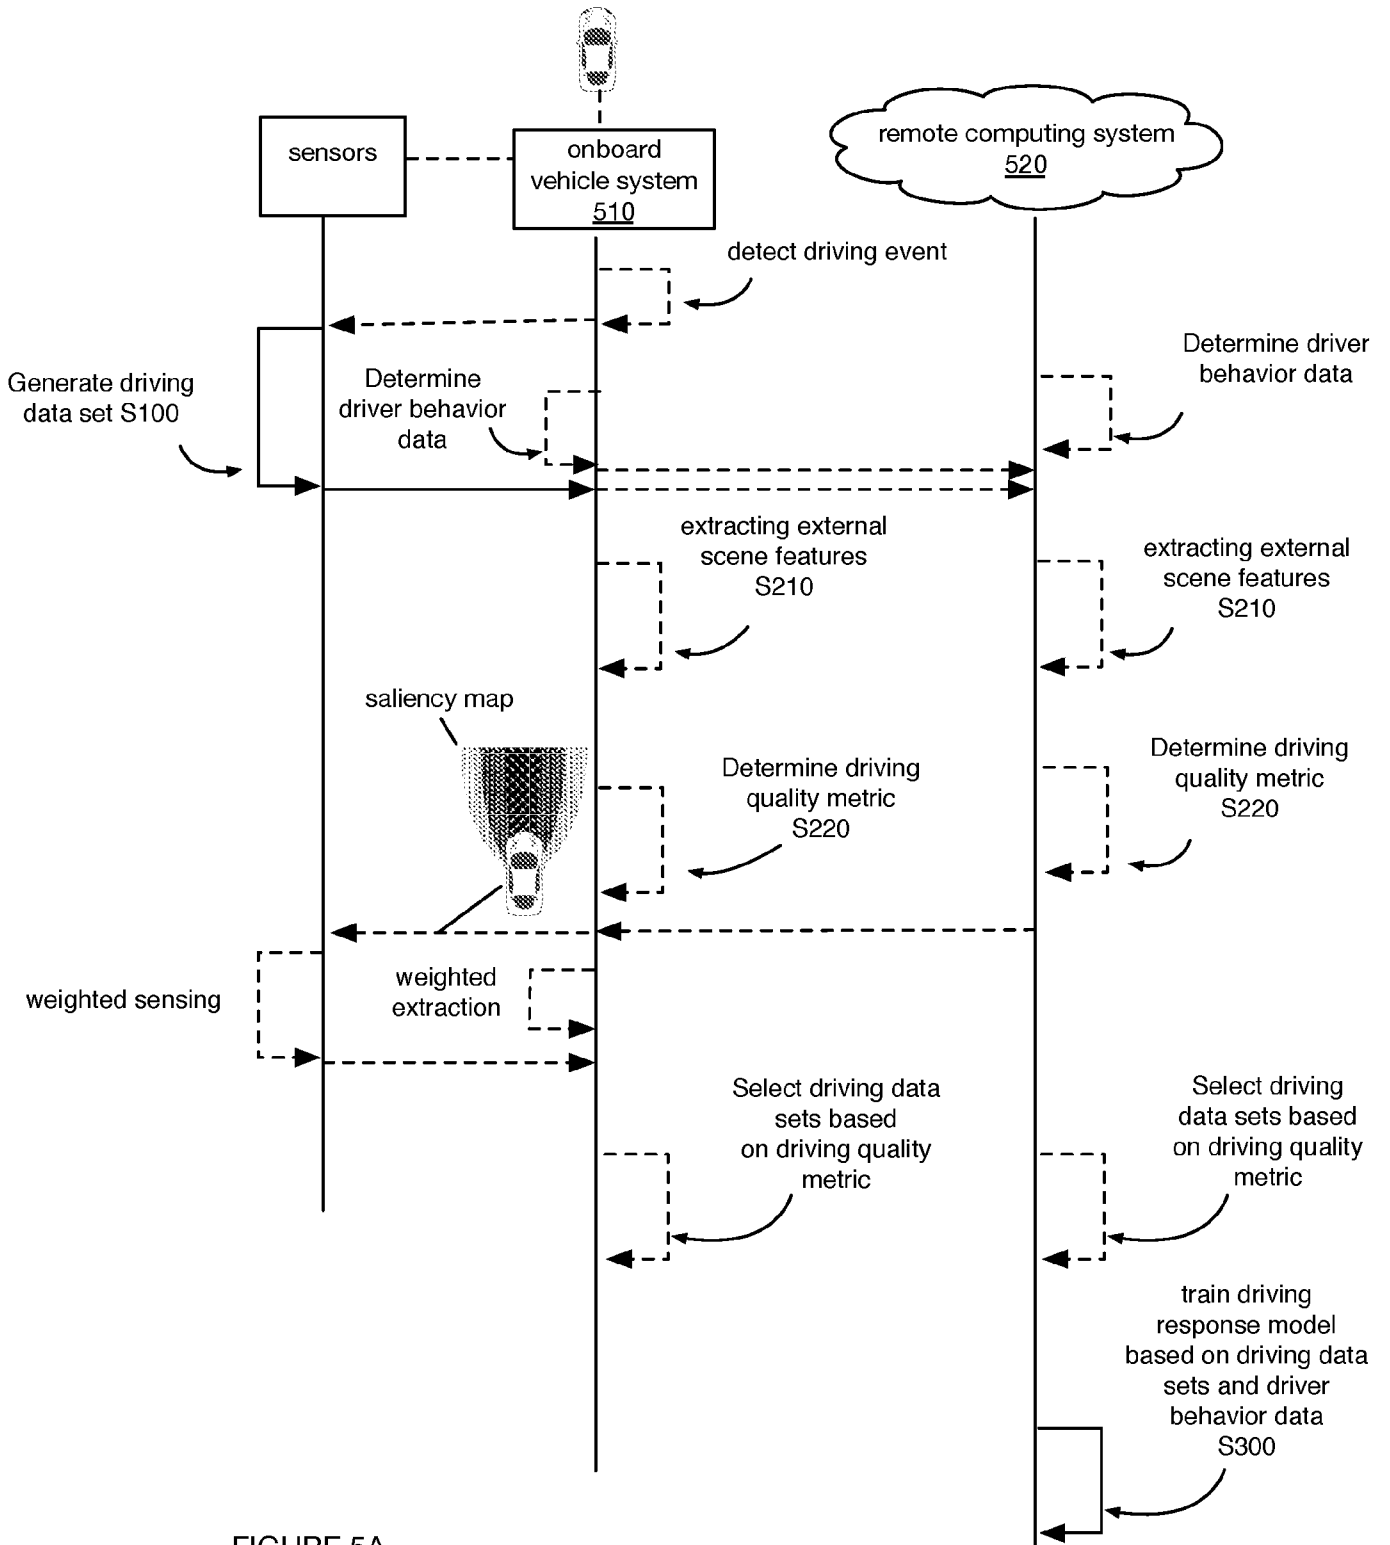
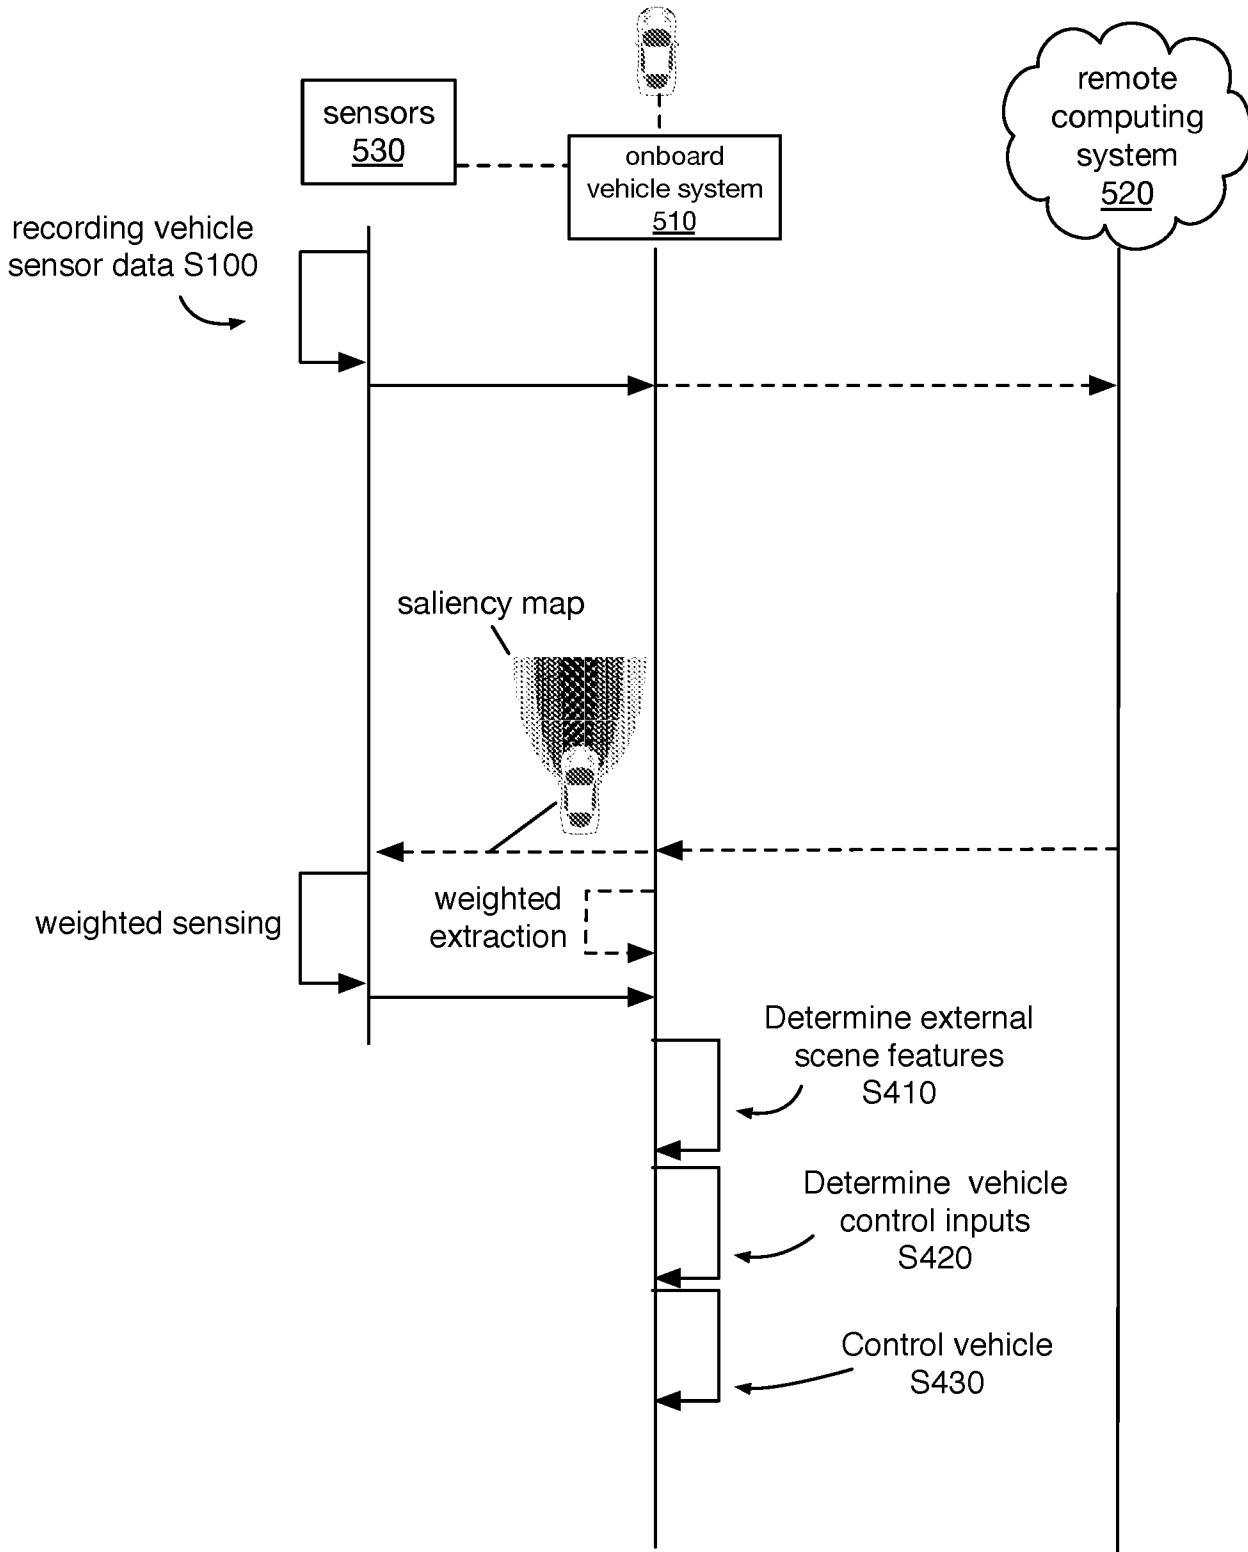
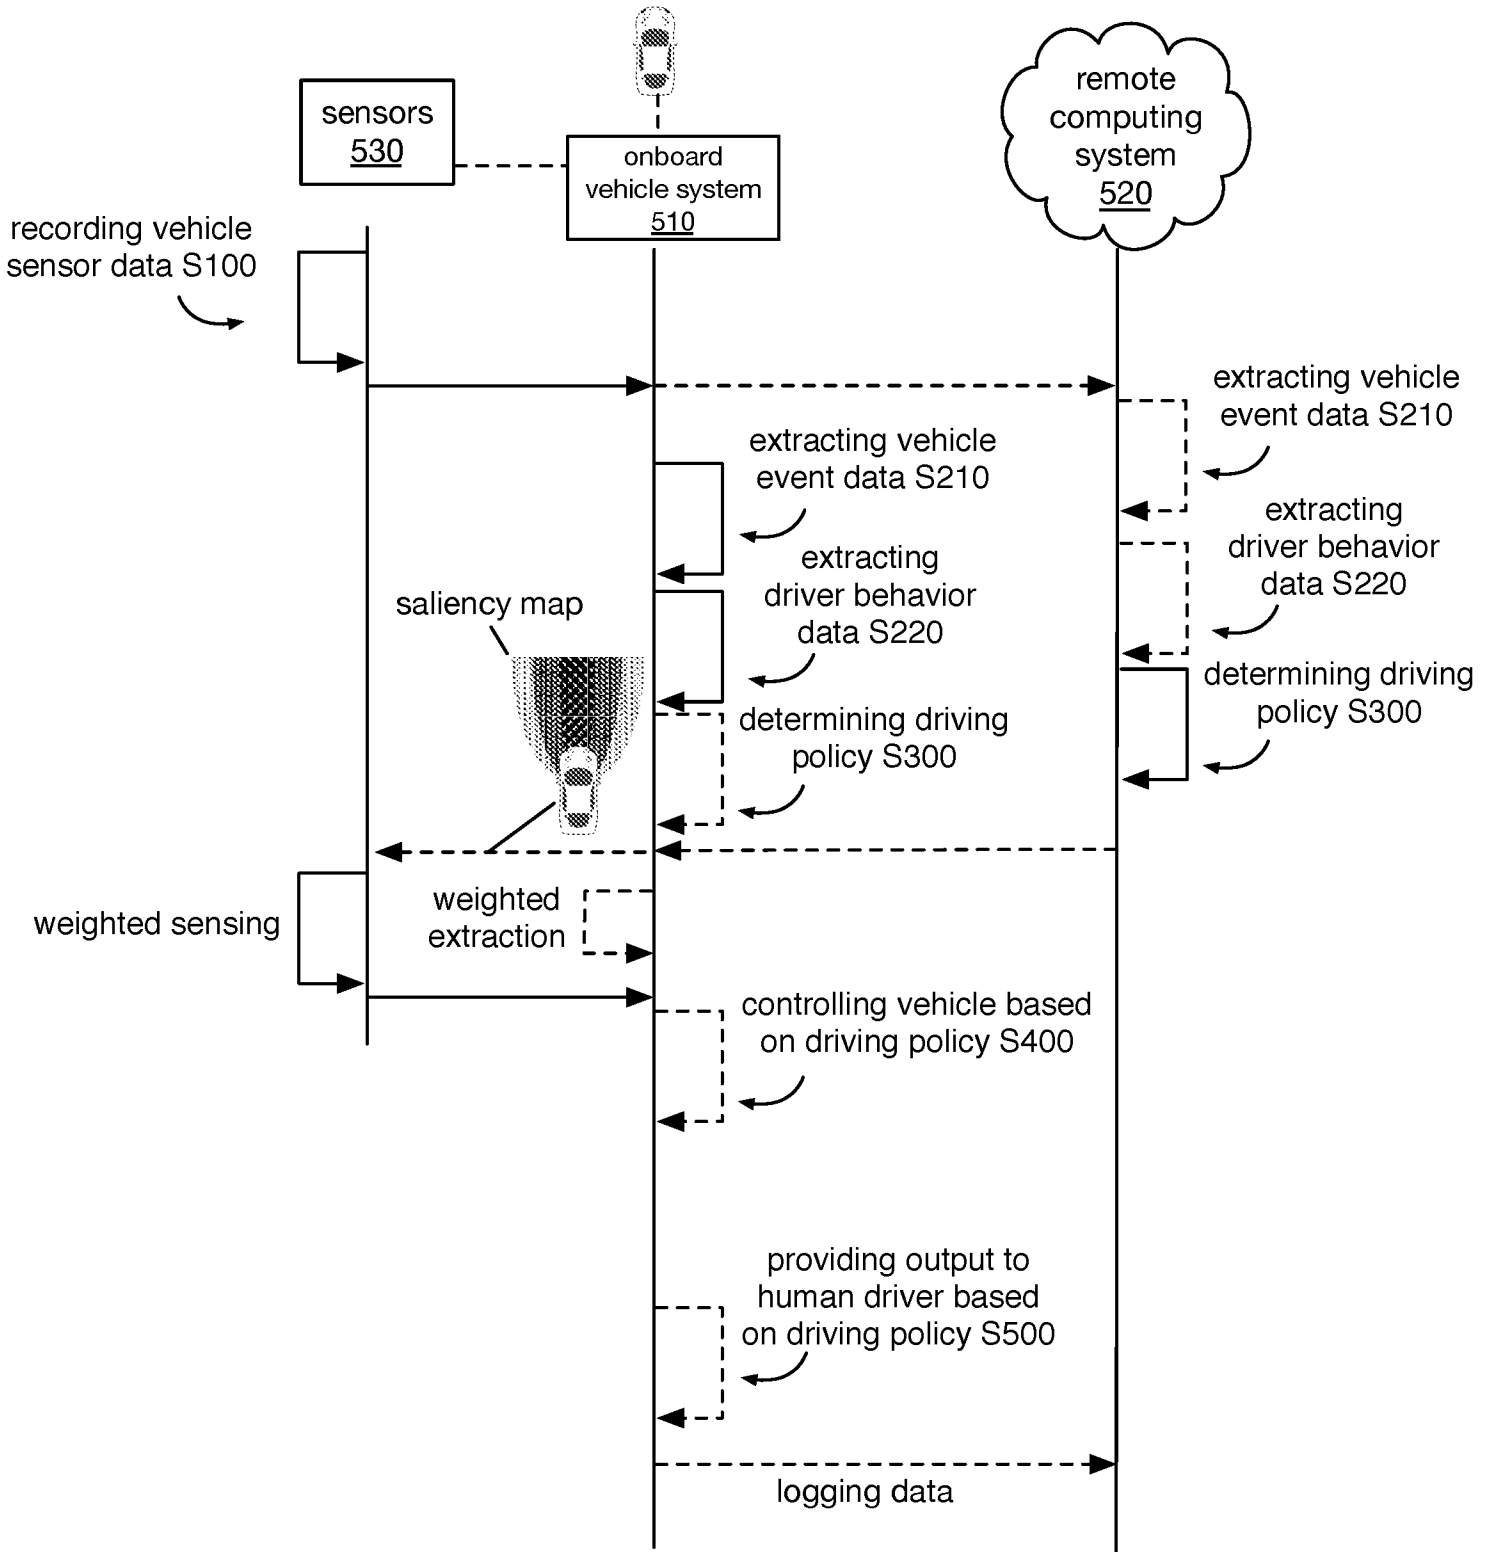


FIGURE 5A

5 / 8



6 / 8



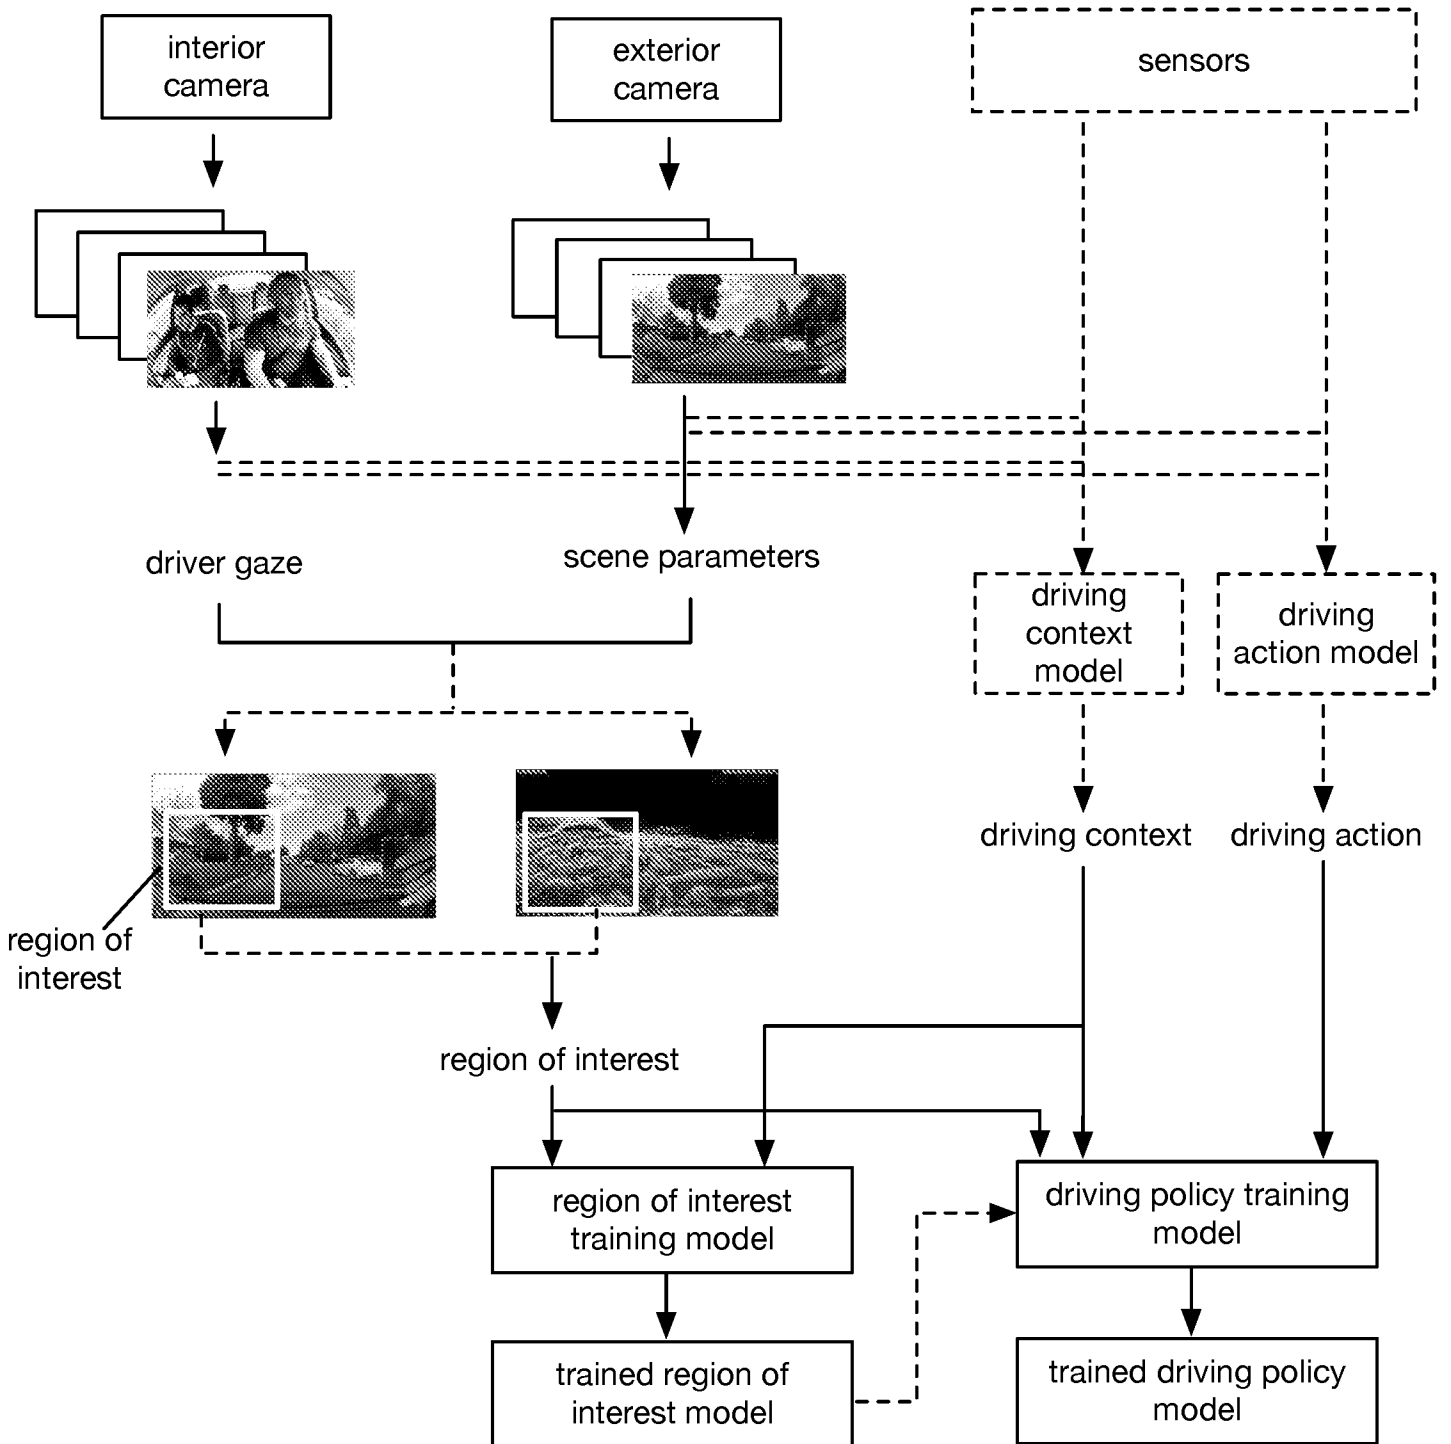


FIGURE 1

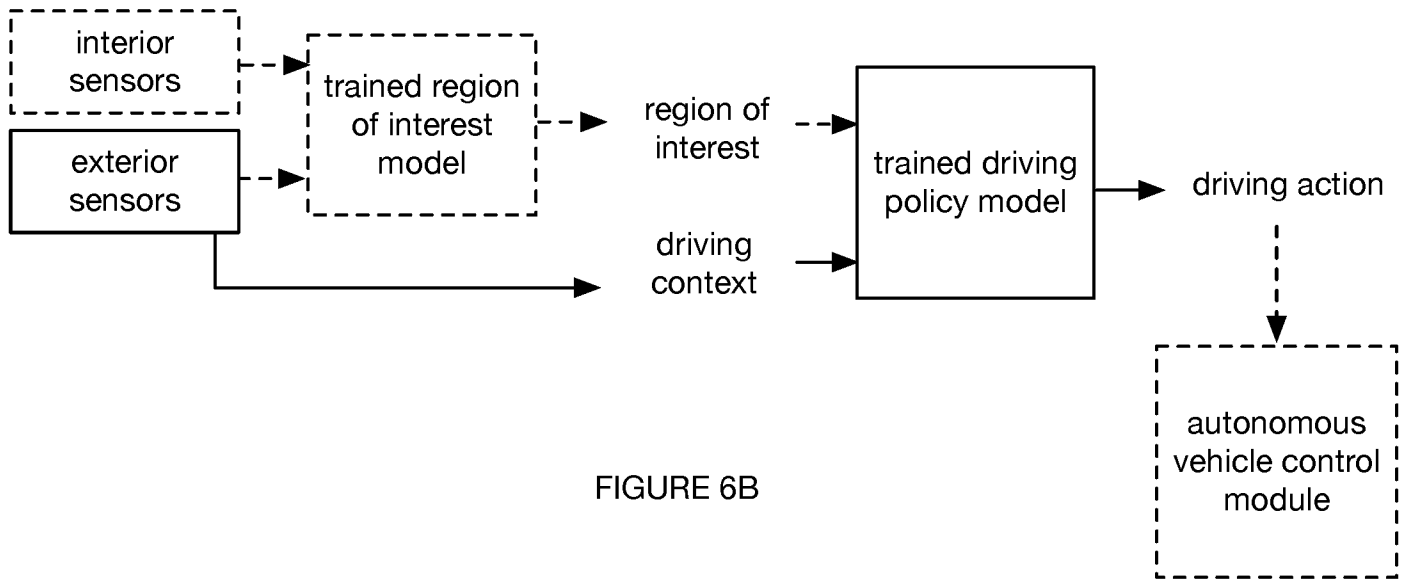


FIGURE 6B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 19/19890

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G08G 1/00 (2019.01)

CPC - G08G 1/166, B60W 40/09, B60K 35/00, G09B 19/167

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History Document

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

See Search History Document

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

See Search History Document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X ---	US 2011/0169625 A1 (JAMES) 14 July 2011 (14.07.2011), entire document, especially para [0006]; [0007]; [0011]; [0014]; [0021]; [0022]; [0029]; [0035]; [0036]; [0042]; [0047]; [0048]; [0054]; [0056]; [0057]; [0059]; [0063]; [0068]; [0076]; [0091]	1-10, 13-20 ----- 11-12
Y	US 2017/0249095 A1 (AUTOCONNECT HOLDINGS, LLC) 31 August 2017 (31.08.2017), entire document, especially para [0112], [0130]; [0382];	11-12

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

24 April 2019 (24.04.2019)

Date of mailing of the international search report

16 MAY 2019

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774