



(43) International Publication Date
20 February 2020 (20.02.2020)

(51) International Patent Classification:
Not classified

(21) International Application Number:
PCT/IB2019/000726

(22) International Filing Date:
14 August 2019 (14.08.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/718,554 14 August 2018 (14.08.2018) US
62/724,355 29 August 2018 (29.08.2018) US
62/772,366 28 November 2018 (28.11.2018) US
62/777,914 11 December 2018 (11.12.2018) US

(71) Applicant: **MOBILEYE VISION TECHNOLOGIES LTD.** [IL/IL]: 13 Hartom Street, P.O.B. 45157, Har Hotzvim, 9777513 Jerusalem (IL).

(72) Inventors: **SHALEV-SHWARTZ, Shai**; c/o Mobileye Vision Technologies Ltd., 13 Hartom Street, P.O.B. 45157, Har Hotzvim, 9777513 Jerusalem (IL). **SHAMMAH, Shaked**; c/o Mobileye Vision Technologies Ltd., 13 Har-

tom Street, P.O.B. 45157, Har Hotzvim, 9777513 Jerusalem (IL). **SHASHUA, Amnon**; c/o Mobileye Vision Technologies Ltd., 13 Hartom Street, P.O.B. 45157, Har Hotzvim, 9777513 Jerusalem (IL). **COHEN, Barak**; c/o Mobileye Vision Technologies Ltd., 13 Hartom Street, P.O.B. 45157, Har Hotzvim, 9777513 Jerusalem (IL). **ADELMAN, Zeev**; c/o Mobileye Vision Technologies Ltd., 13 Hartom Street, P.O.B. 45157, Har Hotzvim, 9777513 Jerusalem (IL).

(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,

(54) Title: SYSTEMS AND METHODS FOR NAVIGATING WITH SAFE DISTANCES





FIG. 49

(57) Abstract: Systems and methods are disclosed for navigating a host vehicle. In one implementation, at least one processing device may be programmed to receive an image representative of an environment of the host vehicle, determine a planned navigational action for the host vehicle, analyze the image to identify a target vehicle with a direction of travel toward the host vehicle, and determine a next-state distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken. The at least one processing device may further determine a stopping distance for the host vehicle based on a braking rate, a maximum acceleration capability, and a current speed of the host vehicle, determine a stopping distance for the target vehicle based on a braking rate, a maximum acceleration capability, and a current speed of the target vehicle, and implement the planned navigational action if the determined next-state distance is greater than a sum of the stopping distances for the host vehicle and the target vehicle.



WO 2020/035728 A2

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).

Published:

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

SYSTEMS AND METHODS FOR NAVIGATING WITH SAFE DISTANCESCross References to Related Applications

[0001] This application claims the benefit of priority of United States Provisional Patent Application No. 62/718,554, filed on August 14, 2018; United States Provisional Patent Application No. 62/724,355, filed on August 29, 2018; United States Provisional Patent Application No. 62/772,366, filed on November 28, 2018; and United States Provisional Patent Application No. 62/777,914, filed on December 11, 2018. All of the foregoing applications are incorporated herein by reference in their entirety.

BACKGROUND10 Technical Field

[0002] The present disclosure relates generally to autonomous vehicle navigation. Additionally, this disclosure relates to systems and methods for navigating according to potential accident liability constraints.

Background Information

15 [0003] As technology continues to advance, the goal of a fully autonomous vehicle that is capable of navigating on roadways is on the horizon. Autonomous vehicles may need to take into account a variety of factors and make appropriate decisions based on those factors to safely and accurately reach an intended destination. For example, an autonomous vehicle may need to process and interpret visual information (e.g., information captured from a camera), information from radar or lidar, and may also use
20 information obtained from other sources (e.g., from a GPS device, a speed sensor, an accelerometer, a suspension sensor, etc.). *At the same time, in order to navigate to a destination, an autonomous vehicle may also need to identify its location within a particular roadway (e.g., a specific lane within a multi-lane road), navigate alongside other vehicles, avoid obstacles and pedestrians, observe traffic signals and signs, travel from one road to another road at appropriate intersections or interchanges, and respond to*
25 *any other situation that occurs or develops during the vehicle's operation. Moreover, the navigational system may need to adhere to certain imposed constraints. In some cases, those constraints may relate to interactions between a host vehicle and one or more other objects, such as other vehicles, pedestrians, etc. In other cases, the constraints may relate to liability rules to be followed in implementing one or more navigational actions for a host vehicle.*

30 [0004] In the field of autonomous driving, there are two important considerations for viable autonomous vehicle systems. The first is a standardization of safety assurance, including requirements that every self-driving car must satisfy to ensure safety, and how those requirements can be verified. The second is scalability, as engineering solutions that lead to unleashed costs will not scale to millions of cars and may prevent widespread or even not so widespread adoption of autonomous vehicles. Thus, there is a
35 need for an interpretable, mathematical model for safety assurance and a design of a system that adheres to safety assurance requirements while being scalable to millions of cars.

SUMMARY

[0005] Embodiments consistent with the present disclosure provide systems and methods for autonomous vehicle navigation. The disclosed embodiments may use cameras to provide autonomous vehicle navigation features. For example, consistent with the disclosed embodiments, the disclosed systems may include one, two, or more cameras that monitor the environment of a vehicle. The disclosed systems may provide a navigational response based on, for example, an analysis of images captured by one or more of the cameras. The navigational response may also take into account other data including, for example, global positioning system (GPS) data, sensor data (e.g., from an accelerometer, a speed sensor, a suspension sensor, etc.), and/or other map data.

[0006] In an embodiment, a system for navigating a host vehicle is disclosed. The system may comprise at least one processing device programmed to receive, from an image capture device, at least one image representative of an environment of the host vehicle; determine, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle; analyze the at least one image to identify a target vehicle in the environment of the host vehicle, wherein a direction of travel of the target vehicle is toward the host vehicle; determine a next-state distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken; determine a host vehicle braking rate, a host vehicle maximum acceleration capability, and a current speed of the host vehicle; determine a stopping distance for the host vehicle based on the host vehicle braking rate, the host vehicle maximum acceleration capability, and the current speed of the host vehicle; determine a current speed of the target vehicle, a target vehicle maximum acceleration capability, and a target vehicle braking rate; determine a stopping distance for the target vehicle based on the target vehicle braking rate, the target vehicle maximum acceleration capability, and the current speed of the target vehicle; and implement the planned navigational action if the determined next-state distance is greater than a sum of the stopping distance for the host vehicle and the stopping distance of the target vehicle.

[0007] In an embodiment, a method for navigating a host vehicle is disclosed. The method may comprise receiving, from an image capture device, at least one image representative of an environment of the host vehicle; determining, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle; analyzing the at least one image to identify a target vehicle in the environment of the host vehicle, wherein a direction of travel of the target vehicle is toward the host vehicle; determining a next-state distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken; determining a host vehicle braking rate, a host vehicle maximum acceleration capability, and a current speed of the host vehicle; determining a stopping distance for the host vehicle based on the host vehicle braking rate, the host vehicle maximum acceleration capability, and the current speed of the host vehicle; determining a current speed of the target vehicle, a target vehicle maximum acceleration capability, and a target vehicle braking rate; determining a stopping distance for the target vehicle based on the target vehicle braking rate, the target vehicle maximum acceleration capability, and the current speed of the target vehicle; and

implementing the planned navigational action if the determined next-state distance is greater than a sum of the stopping distance for the host vehicle and the stopping distance of the target vehicle.

[0008] In an embodiment, a system for navigating a host vehicle is disclosed. The system may comprise at least one processing device programmed to receive, from an image capture device, at least one image representative of an environment of the host vehicle; determine, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle; analyze the at least one image to identify a target vehicle in the environment of the host vehicle; determine a next-state lateral distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken; determine a maximum yaw rate capability of the host vehicle, a maximum change in turn radius capability of the host vehicle, and a current lateral speed of the host vehicle; determine a lateral braking distance for the host vehicle based on the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, and the current lateral speed of the host vehicle; determine a current lateral speed of the target vehicle, a target vehicle maximum yaw rate capability, and a target vehicle maximum change in turn radius capability; determine a lateral braking distance for the target vehicle based on the current lateral speed of the target vehicle, the target vehicle maximum yaw rate capability, and the target vehicle maximum change in turn radius capability; and implement the planned navigational action if the determined next-state lateral distance is greater than a sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle.

[0009] In an embodiment, a method for navigating a host vehicle is disclosed. The method may comprise receiving, from an image capture device, at least one image representative of an environment of the host vehicle; determining, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle; analyzing the at least one image to identify a target vehicle in the environment of the host vehicle; determining a next-state lateral distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken; determining a maximum yaw rate capability of the host vehicle, a maximum change in turn radius capability of the host vehicle, and a current lateral speed of the host vehicle; determining a lateral braking distance for the host vehicle based on the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, and the current lateral speed of the host vehicle; determining a current lateral speed of the target vehicle, a target vehicle maximum yaw rate capability, and a target vehicle maximum change in turn radius capability; determining a lateral braking distance for the target vehicle based on the current lateral speed of the target vehicle, the target vehicle maximum yaw rate capability, and the target vehicle maximum change in turn radius capability; and implementing the planned navigational action if the determined next-state lateral distance is greater than a sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle.

[0010] In an embodiment, a system for navigating a host vehicle in proximity to a pedestrian crosswalk is disclosed. The system may comprise at least one processing device programmed to: receive, from an image capture device, at least one image representative of an environment of the host vehicle;

detect, based on analysis of the at least one image, a representation of a pedestrian crosswalk in the at least one image; determine, based on analysis of the at least one image, whether a representation of a pedestrian appears in the at least one image; detect a presence of a traffic light in the environment of the host vehicle; determine whether the detected traffic light is relevant to the host vehicle and the pedestrian crosswalk; determine a state of the detected traffic light; determine, when a representation of a pedestrian appears in the at least one image, a proximity of the pedestrian relative to the detected pedestrian crosswalk; determine, based on at least one driving policy, a planned navigational action for causing the host vehicle to navigate relative to the detected pedestrian crosswalk, wherein determination of the planned navigational action is further based on the determined state of the detected traffic light and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk; and cause one or more actuator systems of the host vehicle to implement the planned navigational action.

[0011] In an embodiment, a method for navigating a host vehicle in proximity to a pedestrian crosswalk is disclosed. The method may comprise receiving, from an image capture device, at least one image representative of an environment of the host vehicle; detecting, based on analysis of the at least one image, a representation of a pedestrian crosswalk in the at least one image; determining, based on analysis of the at least one image, whether a representation of a pedestrian appears in the at least one image; detecting a presence of a traffic light in the environment of the host vehicle; determining whether the detected traffic light is relevant to the host vehicle and the pedestrian crosswalk; determining a state of the detected traffic light; determining, when a representation of a pedestrian appears in the at least one image, a proximity of the pedestrian relative to the detected pedestrian crosswalk; determining, based on at least one driving policy, a planned navigational action for causing the host vehicle to navigate relative to the detected pedestrian crosswalk, wherein determination of the planned navigational action is further based on the determined state of the detected traffic light and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk; and causing one or more actuator systems of the host vehicle to implement the planned navigational action.

[0012] In an embodiment, a method for navigating a host vehicle in proximity to a pedestrian crosswalk is disclosed. The method may comprise receiving, from an image capture device, at least one image representing an environment of the host vehicle; detecting a starting and an ending location of a crosswalk; determining, based on an analysis of the at least one image, whether a pedestrian is present in a proximity of the crosswalk; detecting a presence of a traffic light in the environment of the host vehicle; determining whether the traffic light is relevant to the host vehicle and the crosswalk; determining a state of the traffic light; determining a navigational action for the host vehicle in vicinity of the detected crosswalk, based on: the relevancy of the traffic light; the determined state of the traffic light; a presence of the pedestrian in the proximity of the crosswalk; a shortest distance selected between the pedestrian and one of the starting or the ending location of the crosswalk; and a motion vector of the pedestrian; and causing one or more actuator systems of the host vehicle to implement the navigational action.

[0013] Consistent with other disclosed embodiments, non-transitory computer-readable storage media may store program instructions, which are executable by at least one processing device and perform any of the steps and/or methods described herein.

[0014] The foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The accompanying drawings, which are incorporated in and constitute a part of this disclosure, illustrate various disclosed embodiments. In the drawings:

[0016] Fig. 1 is a diagrammatic representation of an exemplary system consistent with the disclosed embodiments.

[0017] Fig. 2A is a diagrammatic side view representation of an exemplary vehicle including a system consistent with the disclosed embodiments.

[0018] Fig. 2B is a diagrammatic top view representation of the vehicle and system shown in Fig. 2A consistent with the disclosed embodiments.

[0019] Fig. 2C is a diagrammatic top view representation of another embodiment of a vehicle including a system consistent with the disclosed embodiments.

[0020] Fig. 2D is a diagrammatic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiments.

[0021] Fig. 2E is a diagrammatic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiments.

[0022] Fig. 2F is a diagrammatic representation of exemplary vehicle control systems consistent with the disclosed embodiments.

[0023] Fig. 3A is a diagrammatic representation of an interior of a vehicle including a rearview mirror and a user interface for a vehicle imaging system consistent with the disclosed embodiments.

[0024] Fig. 3B is an illustration of an example of a camera mount that is configured to be positioned behind a rearview mirror and against a vehicle windshield consistent with the disclosed embodiments.

[0025] Fig. 3C is an illustration of the camera mount shown in FIG. 3B from a different perspective consistent with the disclosed embodiments.

[0026] Fig. 3D is an illustration of an example of a camera mount that is configured to be positioned behind a rearview mirror and against a vehicle windshield consistent with the disclosed embodiments.

[0027] Fig. 4 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

[0028] Fig. 5A is a flowchart showing an exemplary process for causing one or more navigational responses based on monocular image analysis consistent with disclosed embodiments.

[0029] Fig. 5B is a flowchart showing an exemplary process for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

[0030] Fig. 5C is a flowchart showing an exemplary process for detecting road marks and/or lane geometry information in a set of images consistent with the disclosed embodiments.

5 [0031] Fig. 5D is a flowchart showing an exemplary process for detecting traffic lights in a set of images consistent with the disclosed embodiments.

[0032] Fig. 5E is a flowchart showing an exemplary process for causing one or more navigational responses based on a vehicle path consistent with the disclosed embodiments.

[0033] Fig. 5F is a flowchart showing an exemplary process for determining whether a leading
10 vehicle is changing lanes consistent with the disclosed embodiments.

[0034] Fig. 6 is a flowchart showing an exemplary process for causing one or more navigational responses based on stereo image analysis consistent with the disclosed embodiments.

[0035] Fig. 7 is a flowchart showing an exemplary process for causing one or more navigational responses based on an analysis of three sets of images consistent with the disclosed
15 embodiments.

[0036] Fig. 8 is a block diagram representation of modules that may be implemented by one or more specifically programmed processing devices of a navigation system for an autonomous vehicle consistent with the disclosed embodiments.

[0037] Fig. 9 is a navigation options graph consistent with the disclosed embodiments.

20 [0038] Fig. 10 is a navigation options graph consistent with the disclosed embodiments.

[0039] Figs. 11A, 11B, and 11C provide a schematic representation of navigational options of a host vehicle in a merge zone consistent with the disclosed embodiments.

[0040] Fig. 11D provide a diagrammatic depiction of a double merge scenario consistent with the disclosed embodiments.

25 [0041] Fig. 11E provides an options graph potentially useful in a double merge scenario consistent with the disclosed embodiments.

[0042] Fig. 12 provides a diagram of a representative image captured of an environment of a host vehicle, along with potential navigational constraints consistent with the disclosed embodiments.

[0043] Fig. 13 provides an algorithmic flow chart for navigating a vehicle consistent with the
30 disclosed embodiments.

[0044] Fig. 14 provides an algorithmic flow chart for navigating a vehicle consistent with the disclosed embodiments.

[0045] Fig. 15 provides an algorithmic flow chart for navigating a vehicle consistent with the disclosed embodiments.

35 [0046] Fig. 16 provides an algorithmic flow chart for navigating a vehicle consistent with the disclosed embodiments.

[0047] Figs. 17A and 17B provide a diagrammatic illustration of a host vehicle navigating into a roundabout consistent with the disclosed embodiments.

[0048] Fig. 18 provides an algorithmic flow chart for navigating a vehicle consistent with the disclosed embodiments.

[0049] Fig. 19 illustrates an example of a host vehicle driving on a multi-lane highway consistent with the disclosed embodiments.

5 [0050] Figs. 20A and 20B illustrate examples of a vehicle cutting in in front of another vehicle consistent with the disclosed embodiments.

[0051] Fig. 21 illustrates an example of a vehicle following another vehicle consistent with the disclosed embodiments.

10 [0052] Fig. 22 illustrates an example of a vehicle exiting a parking lot and merging into a possibly busy road consistent with the disclosed embodiments.

[0053] Fig. 23 illustrates a vehicle traveling on a road consistent with the disclosed embodiments.

[0054] Figs. 24A-24D illustrate four example scenarios consistent with the disclosed embodiments.

15 [0055] Fig. 25 illustrates an example scenario consistent with the disclosed embodiments.

[0056] Fig. 26 illustrates an example scenario consistent with the disclosed embodiments.

[0057] Fig. 27 illustrates an example scenario consistent with the disclosed embodiments.

[0058] Figs. 28A and 28B illustrate an example of a scenario in which a vehicle is following another vehicle consistent with the disclosed embodiments.

20 [0059] Figs. 29A and 29B illustrate example blame in cut-in scenarios consistent with the disclosed embodiments.

[0060] Figs. 30A and 30B illustrate example blame in cut-in scenarios consistent with the disclosed embodiments.

25 [0061] Figs. 31A-31D illustrate example blame in drifting scenarios consistent with the disclosed embodiments.

[0062] Figs. 32A and 32B illustrate example blame in two-way traffic scenarios consistent with the disclosed embodiments.

[0063] Figs. 33A and 33B illustrate example blame in two-way traffic scenarios consistent with the disclosed embodiments.

30 [0064] Figs. 34A and 34B illustrate example blame in route priority scenarios consistent with the disclosed embodiments.

[0065] Figs. 35A and 35B illustrate example blame in route priority scenarios consistent with the disclosed embodiments.

35 [0066] Figs. 36A and 36B illustrate example blame in route priority scenarios consistent with the disclosed embodiments.

[0067] Figs. 37A and 37B illustrate example blame in route priority scenarios consistent with the disclosed embodiments.

[0068] Figs. 38A and 38B illustrate example blame in route priority scenarios consistent with the disclosed embodiments.

[0069] Figs. 39A and 39B illustrate example blame in route priority scenarios consistent with the disclosed embodiments.

5 [0070] Figs. 40A and 40B illustrate example blame in traffic light scenarios consistent with the disclosed embodiments.

[0071] Figs. 41A and 41B illustrate example blame in traffic light scenarios consistent with the disclosed embodiments.

10 [0072] Figs. 42A and 42B illustrate example blame in traffic light scenarios consistent with the disclosed embodiments.

[0073] Figs. 43A-43C illustrate example vulnerable road users (VRUs) scenarios consistent with the disclosed embodiments.

[0074] Figs. 44A-44C illustrate example vulnerable road users (VRUs) scenarios consistent with the disclosed embodiments.

15 [0075] Figs. 45A-45C illustrate example vulnerable road users (VRUs) scenarios consistent with the disclosed embodiments.

[0076] Figs. 46A-46D illustrate example vulnerable road users (VRUs) scenarios consistent with the disclosed embodiments.

20 [0077] FIG. 47A is an illustration of blame time and proper responses, consistent with the disclosed embodiments.

[0078] FIG. 47B is an illustration of route priority for routes of differing geometries, consistent with the disclosed embodiments.

[0079] FIG. 47C is an illustration of longitudinal ordering on routes of differing geometries, consistent with the disclosed embodiments.

25 [0080] FIG. 47D is an illustration of safe longitudinal distance between vehicles, consistent with the disclosed embodiments.

[0081] FIG. 47E is an illustration of a situation in which a vehicle cannot predict the path of another vehicle, consistent with the disclosed embodiments.

30 [0082] FIG. 47F is an illustration of route priority at a traffic light, consistent with the disclosed embodiments.

[0083] FIG. 47G is an illustration of exemplary unstructured routes, consistent with the disclosed embodiments.

[0084] FIG. 47H is an illustration of exemplary lateral behavior in an unstructured situation, consistent with the disclosed embodiments.

35 [0085] FIG. 47I is an illustration of exposure time and blame time at an area of occlusion, consistent with the disclosed embodiments.

[0086] Fig. 48A illustrates an example scenario of two vehicles traveling in opposite directions consistent with the disclosed embodiments.

[0087] Fig. 48B illustrates an example of a target vehicle traveling towards a host vehicle consistent with the disclosed embodiments.

[0088] Fig. 49 illustrates an example of host vehicle maintaining a safe longitudinal distance consistent with the disclosed embodiments.

5 [0089] Figs. 50A and 50B provide a flowchart depicting an exemplary process for maintaining a safe longitudinal distance consistent with the disclosed embodiments.

[0090] Fig. 51A illustrates an example scenario with two vehicles spaced laterally from each other consistent with the disclosed embodiments.

10 [0091] Fig. 51B illustrates an example of host vehicle maintaining a safe lateral distance consistent with the disclosed embodiments.

[0092] Fig. 52A illustrates an example of a host vehicle performing a planned navigation action consistent with the disclosed embodiments.

[0093] Fig. 52B shows an example of a host vehicle determining whether to perform a navigation action.

15 [0094] Figs. 53A and 53B provide a flowchart depicting an example process for maintaining a safe lateral distance consistent with the disclosed embodiments.

[0095] Fig. 54 is a schematic illustration of a roadway including crosswalks, consistent with the disclosed embodiments.

20 [0096] Figs. 55A and 55B are schematic diagrams of possible navigational actions executed by vehicles traveling along a roadway, consistent with the disclosed embodiments.

[0097] Figs. 55C and 55D show examples of estimating a distance from a pedestrian to a crosswalk, consistent with the disclosed embodiments.

[0098] Fig. 56 is a flowchart describing a process of navigating a vehicle in the proximity of a crosswalk, consistent with the disclosed embodiments.

25 **DETAILED DESCRIPTION**

[0099] The following detailed description refers to the accompanying drawings. Wherever possible, the same reference numbers are used in the drawings and the following description to refer to the same or similar parts. While several illustrative embodiments are described herein, modifications, adaptations and other implementations are possible. For example, substitutions, additions or
30 modifications may be made to the components illustrated in the drawings, and the illustrative methods described herein may be modified by substituting, reordering, removing, or adding steps to the disclosed methods. Accordingly, the following detailed description is not limited to the disclosed embodiments and examples. Instead, the proper scope is defined by the appended claims.

[0100] Autonomous Vehicle Overview

35 [0101] As used throughout this disclosure, the term “autonomous vehicle” refers to a vehicle capable of implementing at least one navigational change without driver input. A “navigational change” refers to a change in one or more of steering, braking, or acceleration/deceleration of the vehicle. To be autonomous, a vehicle need not be fully automatic (e.g., fully operational without a driver or without

driver input). Rather, an autonomous vehicle includes those that can operate under driver control during certain time periods and without driver control during other time periods. Autonomous vehicles may also include vehicles that control only some aspects of vehicle navigation, such as steering (e.g., to maintain a vehicle course between vehicle lane constraints) or some steering operations under certain circumstances (but not under all circumstances), but may leave other aspects to the driver (e.g., braking or braking under certain circumstances). In some cases, autonomous vehicles may handle some or all aspects of braking, speed control, and/or steering of the vehicle.

[0102] As human drivers typically rely on visual cues and observations in order to control a vehicle, transportation infrastructures are built accordingly, with lane markings, traffic signs, and traffic lights designed to provide visual information to drivers. In view of these design characteristics of transportation infrastructures, an autonomous vehicle may include a camera and a processing unit that analyzes visual information captured from the environment of the vehicle. The visual information may include, for example, images representing components of the transportation infrastructure (e.g., lane markings, traffic signs, traffic lights, etc.) that are observable by drivers and other obstacles (e.g., other vehicles, pedestrians, debris, etc.). Additionally, an autonomous vehicle may also use stored information, such as information that provides a model of the vehicle's environment when navigating. For example, the vehicle may use GPS data, sensor data (e.g., from an accelerometer, a speed sensor, a suspension sensor, etc.), and/or other map data to provide information related to its environment while it is traveling, and the vehicle (as well as other vehicles) may use the information to localize itself on the model. Some vehicles can also be capable of communication among them, sharing information, altering the peer vehicle of hazards or changes in the vehicles' surroundings, etc.

[0103] System Overview

[0104] Fig. 1 is a block diagram representation of a system 100 consistent with the exemplary disclosed embodiments. System 100 may include various components depending on the requirements of a particular implementation. In some embodiments, system 100 may include a processing unit 110, an image acquisition unit 120, a position sensor 130, one or more memory units 140, 150, a map database 160, a user interface 170, and a wireless transceiver 172. Processing unit 110 may include one or more processing devices. In some embodiments, processing unit 110 may include an applications processor 180, an image processor 190, or any other suitable processing device. Similarly, image acquisition unit 120 may include any number of image acquisition devices and components depending on the requirements of a particular application. In some embodiments, image acquisition unit 120 may include one or more image capture devices (e.g., cameras, CCDs, or any other type of image sensor), such as image capture device 122, image capture device 124, and image capture device 126. System 100 may also include a data interface 128 communicatively connecting processing unit 110 to image acquisition unit 120. For example, data interface 128 may include any wired and/or wireless link or links for transmitting image data acquired by image acquisition unit 120 to processing unit 110.

[0105] Wireless transceiver 172 may include one or more devices configured to exchange transmissions over an air interface to one or more networks (e.g., cellular, the Internet, etc.) by use of a

radio frequency, infrared frequency, magnetic field, or an electric field. Wireless transceiver 172 may use any known standard to transmit and/or receive data (e.g., Wi-Fi, Bluetooth®, Bluetooth Smart, 802.15.4, ZigBee, etc.). Such transmissions can include communications from the host vehicle to one or more remotely located servers. Such transmissions may also include communications (one-way or two-way) between the host vehicle and one or more target vehicles in an environment of the host vehicle (e.g., to facilitate coordination of navigation of the host vehicle in view of or together with target vehicles in the environment of the host vehicle), or even a broadcast transmission to unspecified recipients in a vicinity of the transmitting vehicle.

[0106] Both applications processor 180 and image processor 190 may include various types of hardware-based processing devices. For example, either or both of applications processor 180 and image processor 190 may include a microprocessor, preprocessors (such as an image preprocessor), graphics processors, a central processing unit (CPU), support circuits, digital signal processors, integrated circuits, memory, or any other types of devices suitable for running applications and for image processing and analysis. In some embodiments, applications processor 180 and/or image processor 190 may include any type of single or multi-core processor, mobile device microcontroller, central processing unit, etc. Various processing devices may be used, including, for example, processors available from manufacturers such as Intel®, AMD®, etc. and may include various architectures (e.g., x86 processor, ARM®, etc.).

[0107] In some embodiments, applications processor 180 and/or image processor 190 may include any of the EyeQ series of processor chips available from Mobileye®. These processor designs each include multiple processing units with local memory and instruction sets. Such processors may include video inputs for receiving image data from multiple image sensors and may also include video output capabilities. In one example, the EyeQ2® uses 90nm-micron technology operating at 332Mhz. The EyeQ2® architecture consists of two floating point, hyper-thread 32-bit RISC CPUs (MIPS32® 34K® cores), five Vision Computing Engines (VCE), three Vector Microcode Processors (VMP®), Denali 64-bit Mobile DDR Controller, 128-bit internal Sonics Interconnect, dual 16-bit Video input and 18-bit Video output controllers, 16 channels DMA and several peripherals. The MIPS34K CPU manages the five VCEs, three VMP™ and the DMA, the second MIPS34K CPU and the multi-channel DMA as well as the other peripherals. The five VCEs, three VMP® and the MIPS34K CPU can perform intensive vision computations required by multi-function bundle applications. In another example, the EyeQ3®, which is a third generation processor and is six times more powerful than the EyeQ2®, may be used in the disclosed embodiments. In other examples, the EyeQ4® and/or the EyeQ5® may be used in the disclosed embodiments. Of course, any newer or future EyeQ processing devices may also be used together with the disclosed embodiments.

[0108] Any of the processing devices disclosed herein may be configured to perform certain functions. Configuring a processing device, such as any of the described EyeQ processors or other controller or microprocessor, to perform certain functions may include programming of computer executable instructions and making those instructions available to the processing device for execution during operation of the processing device. In some embodiments, configuring a processing device may

include programming the processing device directly with architectural instructions. In other embodiments, configuring a processing device may include storing executable instructions on a memory that is accessible to the processing device during operation. For example, the processing device may access the memory to obtain and execute the stored instructions during operation. In either case, the processing device configured to perform the sensing, image analysis, and/or navigational functions disclosed herein represents a specialized hardware-based system in control of multiple hardware based components of a host vehicle.

[0109] While Fig. 1 depicts two separate processing devices included in processing unit 110, more or fewer processing devices may be used. For example, in some embodiments, a single processing device may be used to accomplish the tasks of applications processor 180 and image processor 190. In other embodiments, these tasks may be performed by more than two processing devices. Further, in some embodiments, system 100 may include one or more of processing unit 110 without including other components, such as image acquisition unit 120.

[0110] Processing unit 110 may comprise various types of devices. For example, processing unit 110 may include various devices, such as a controller, an image preprocessor, a central processing unit (CPU), support circuits, digital signal processors, integrated circuits, memory, or any other types of devices for image processing and analysis. The image preprocessor may include a video processor for capturing, digitizing and processing the imagery from the image sensors. The CPU may comprise any number of microcontrollers or microprocessors. The support circuits may be any number of circuits generally well known in the art, including cache, power supply, clock and input-output circuits. The memory may store software that, when executed by the processor, controls the operation of the system. The memory may include databases and image processing software. The memory may comprise any number of random access memories, read only memories, flash memories, disk drives, optical storage, tape storage, removable storage and other types of storage. In one instance, the memory may be separate from the processing unit 110. In another instance, the memory may be integrated into the processing unit 110.

[0111] Each memory 140, 150 may include software instructions that when executed by a processor (e.g., applications processor 180 and/or image processor 190), may control operation of various aspects of system 100. These memory units may include various databases and image processing software, as well as a trained system, such as a neural network, or a deep neural network, for example. The memory units may include random access memory, read only memory, flash memory, disk drives, optical storage, tape storage, removable storage and/or any other types of storage. In some embodiments, memory units 140, 150 may be separate from the applications processor 180 and/or image processor 190. In other embodiments, these memory units may be integrated into applications processor 180 and/or image processor 190.

[0112] Position sensor 130 may include any type of device suitable for determining a location associated with at least one component of system 100. In some embodiments, position sensor 130 may include a GPS receiver. Such receivers can determine a user position and velocity by processing signals

broadcasted by global positioning system satellites. Position information from position sensor 130 may be made available to applications processor 180 and/or image processor 190.

[0113] In some embodiments, system 100 may include components such as a speed sensor (e.g., a speedometer) for measuring a speed of vehicle 200. System 100 may also include one or more accelerometers (either single axis or multiaxis) for measuring accelerations of vehicle 200 along one or more axes.

[0114] The memory units 140, 150 may include a database, or data organized in any other form, that indication a location of known landmarks. Sensory information (such as images, radar signal, depth information from lidar or stereo processing of two or more images) of the environment may be processed together with position information, such as a GPS coordinate, vehicle's ego motion, etc. to determine a current location of the vehicle relative to the known landmarks, and refine the vehicle location. Certain aspects of this technology are included in a localization technology known as REMTM, which is being marketed by the assignee of the present application.

[0115] User interface 170 may include any device suitable for providing information to or for receiving inputs from one or more users of system 100. In some embodiments, user interface 170 may include user input devices, including, for example, a touchscreen, microphone, keyboard, pointer devices, track wheels, cameras, knobs, buttons, etc. With such input devices, a user may be able to provide information inputs or commands to system 100 by typing instructions or information, providing voice commands, selecting menu options on a screen using buttons, pointers, or eye-tracking capabilities, or through any other suitable techniques for communicating information to system 100.

[0116] User interface 170 may be equipped with one or more processing devices configured to provide and receive information to or from a user and process that information for use by, for example, applications processor 180. In some embodiments, such processing devices may execute instructions for recognizing and tracking eye movements, receiving and interpreting voice commands, recognizing and interpreting touches and/or gestures made on a touchscreen, responding to keyboard entries or menu selections, etc. In some embodiments, user interface 170 may include a display, speaker, tactile device, and/or any other devices for providing output information to a user.

[0117] Map database 160 may include any type of database for storing map data useful to system 100. In some embodiments, map database 160 may include data relating to the position, in a reference coordinate system, of various items, including roads, water features, geographic features, businesses, points of interest, restaurants, gas stations, etc. Map database 160 may store not only the locations of such items, but also descriptors relating to those items, including, for example, names associated with any of the stored features. In some embodiments, map database 160 may be physically located with other components of system 100. Alternatively or additionally, map database 160 or a portion thereof may be located remotely with respect to other components of system 100 (e.g., processing unit 110). In such embodiments, information from map database 160 may be downloaded over a wired or wireless data connection to a network (e.g., over a cellular network and/or the Internet, etc.). In some cases, map database 160 may store a sparse data model including polynomial representations of certain

road features (e.g., lane markings) or target trajectories for the host vehicle. Map database 160 may also include stored representations of various recognized landmarks that may be used to determine or update a known position of the host vehicle with respect to a target trajectory. The landmark representations may include data fields such as landmark type, landmark location, among other potential identifiers.

5 [0118] Image capture devices 122, 124, and 126 may each include any type of device suitable for capturing at least one image from an environment. Moreover, any number of image capture devices may be used to acquire images for input to the image processor. Some embodiments may include only a single image capture device, while other embodiments may include two, three, or even four or more image capture devices. Image capture devices 122, 124, and 126 will be further described with reference
10 to Figs. 2B-2E, below.

[0119] One or more cameras (e.g., image capture devices 122, 124, and 126) may be part of a sensing block included on a vehicle. Various other sensors may be included in the sensing block, and any or all of the sensors may be relied upon to develop a sensed navigational state of the vehicle. In addition to cameras (forward, sideward, rearward, etc), other sensors such as RADAR, LIDAR, and acoustic
15 sensors may be included in the sensing block. Additionally, the sensing block may include one or more components configured to communicate and transmit/receive information relating to the environment of the vehicle. For example, such components may include wireless transceivers (RF, etc.) that may receive from a source remotely located with respect to the host vehicle sensor based information or any other type of information relating to the environment of the host vehicle. Such information may include sensor
20 output information, or related information, received from vehicle systems other than the host vehicle. In some embodiments, such information may include information received from a remote computing device, a centralized server, etc. Furthermore, the cameras may take on many different configurations: single camera units, multiple cameras, camera clusters, long FOV, short FOV, wide angle, fisheye, etc.

[0120] System 100, or various components thereof, may be incorporated into various different
25 platforms. In some embodiments, system 100 may be included on a vehicle 200, as shown in Fig. 2A. For example, vehicle 200 may be equipped with a processing unit 110 and any of the other components of system 100, as described above relative to FIG. 1. While in some embodiments vehicle 200 may be equipped with only a single image capture device (e.g., camera), in other embodiments, such as those discussed in connection with FIGS. 2B-2E, multiple image capture devices may be used. For example,
30 either of image capture devices 122 and 124 of vehicle 200, as shown in Fig. 2A, may be part of an ADAS (Advanced Driver Assistance Systems) imaging set.

[0121] The image capture devices included on vehicle 200 as part of the image acquisition unit
120 may be positioned at any suitable location. In some embodiments, as shown in Figs. 2A-2E and 3A-3C, image capture device 122 may be located in the vicinity of the rearview mirror. This position may
35 provide a line of sight similar to that of the driver of vehicle 200, which may aid in determining what is and is not visible to the driver. Image capture device 122 may be positioned at any location near the rearview mirror, but placing image capture device 122 on the driver side of the mirror may further aid in obtaining images representative of the driver's field of view and/or line of sight.

[0122] Other locations for the image capture devices of image acquisition unit 120 may also be used. For example, image capture device 124 may be located on or in a bumper of vehicle 200. Such a location may be especially suitable for image capture devices having a wide field of view. The line of sight of bumper-located image capture devices can be different from that of the driver and, therefore, the bumper image capture device and driver may not always see the same objects. The image capture devices (e.g., image capture devices 122, 124, and 126) may also be located in other locations. For example, the image capture devices may be located on or in one or both of the side mirrors of vehicle 200, on the roof of vehicle 200, on the hood of vehicle 200, on the trunk of vehicle 200, on the sides of vehicle 200, mounted on, positioned behind, or positioned in front of any of the windows of vehicle 200, and mounted in or near light fixtures on the front and/or back of vehicle 200, etc.

[0123] In addition to image capture devices, vehicle 200 may include various other components of system 100. For example, processing unit 110 may be included on vehicle 200 either integrated with or separate from an engine control unit (ECU) of the vehicle. Vehicle 200 may also be equipped with a position sensor 130, such as a GPS receiver and may also include a map database 160 and memory units 140 and 150.

[0124] As discussed earlier, wireless transceiver 172 may and/or receive data over one or more networks (e.g., cellular networks, the Internet, etc.). For example, wireless transceiver 172 may upload data collected by system 100 to one or more servers, and download data from the one or more servers. Via wireless transceiver 172, system 100 may receive, for example, periodic or on demand updates to data stored in map database 160, memory 140, and/or memory 150. Similarly, wireless transceiver 172 may upload any data (e.g., images captured by image acquisition unit 120, data received by position sensor 130 or other sensors, vehicle control systems, etc.) from system 100 and/or any data processed by processing unit 110 to the one or more servers.

[0125] System 100 may upload data to a server (e.g., to the cloud) based on a privacy level setting. For example, system 100 may implement privacy level settings to regulate or limit the types of data (including metadata) sent to the server that may uniquely identify a vehicle and or driver/owner of a vehicle. Such settings may be set by user via, for example, wireless transceiver 172, be initialized by factory default settings, or by data received by wireless transceiver 172.

[0126] In some embodiments, system 100 may upload data according to a “high” privacy level, and under setting a setting, system 100 may transmit data (e.g., location information related to a route, captured images, etc.) without any details about the specific vehicle and/or driver/owner. For example, when uploading data according to a “high” privacy setting, system 100 may not include a vehicle identification number (VIN) or a name of a driver or owner of the vehicle, and may instead transmit data, such as captured images and/or limited location information related to a route.

[0127] Other privacy levels are contemplated as well. For example, system 100 may transmit data to a server according to an “intermediate” privacy level and include additional information not included under a “high” privacy level, such as a make and/or model of a vehicle and/or a vehicle type (e.g., a passenger vehicle, sport utility vehicle, truck, etc.). In some embodiments, system 100 may

upload data according to a “low” privacy level. Under a “low” privacy level setting, system 100 may upload data and include information sufficient to uniquely identify a specific vehicle, owner/driver, and/or a portion or entirety of a route traveled by the vehicle. Such “low” privacy level data may include one or more of, for example, a VIN, a driver/owner name, an origination point of a vehicle prior to departure, an intended destination of the vehicle, a make and/or model of the vehicle, a type of the vehicle, etc.

[0128] Fig. 2A is a diagrammatic side view representation of an exemplary vehicle imaging system consistent with the disclosed embodiments. Fig. 2B is a diagrammatic top view illustration of the embodiment shown in Fig. 2A. As illustrated in Fig. 2B, the disclosed embodiments may include a vehicle 200 including in its body a system 100 with a first image capture device 122 positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200, a second image capture device 124 positioned on or in a bumper region (e.g., one of bumper regions 210) of vehicle 200, and a processing unit 110.

[0129] As illustrated in Fig. 2C, image capture devices 122 and 124 may both be positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200. Additionally, while two image capture devices 122 and 124 are shown in Figs. 2B and 2C, it should be understood that other embodiments may include more than two image capture devices. For example, in the embodiments shown in Figs. 2D and 2E, first, second, and third image capture devices 122, 124, and 126, are included in the system 100 of vehicle 200.

[0130] As illustrated in Fig. 2D, image capture device 122 may be positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200, and image capture devices 124 and 126 may be positioned on or in a bumper region (e.g., one of bumper regions 210) of vehicle 200. And as shown in Fig. 2E, image capture devices 122, 124, and 126 may be positioned in the vicinity of the rearview mirror and/or near the driver seat of vehicle 200. The disclosed embodiments are not limited to any particular number and configuration of the image capture devices, and the image capture devices may be positioned in any appropriate location within and/or on vehicle 200.

[0131] It is to be understood that the disclosed embodiments are not limited to vehicles and could be applied in other contexts. It is also to be understood that disclosed embodiments are not limited to a particular type of vehicle 200 and may be applicable to all types of vehicles including automobiles, trucks, trailers, and other types of vehicles.

[0132] The first image capture device 122 may include any suitable type of image capture device. Image capture device 122 may include an optical axis. In one instance, the image capture device 122 may include an Aptina M9V024 WVGA sensor with a global shutter. In other embodiments, image capture device 122 may provide a resolution of 1280x960 pixels and may include a rolling shutter. Image capture device 122 may include various optical elements. In some embodiments one or more lenses may be included, for example, to provide a desired focal length and field of view for the image capture device. In some embodiments, image capture device 122 may be associated with a 6mm lens or a 12mm lens. In some embodiments, image capture device 122 may be configured to capture images having a desired

field-of-view (FOV) 202, as illustrated in Fig. 2D. For example, image capture device 122 may be configured to have a regular FOV, such as within a range of 40 degrees to 56 degrees, including a 46 degree FOV, 50 degree FOV, 52 degree FOV, or greater. Alternatively, image capture device 122 may be configured to have a narrow FOV in the range of 23 to 40 degrees, such as a 28 degree FOV or 36 degree FOV. In addition, image capture device 122 may be configured to have a wide FOV in the range of 100 to 180 degrees. In some embodiments, image capture device 122 may include a wide angle bumper camera or one with up to a 180 degree FOV. In some embodiments, image capture device 122 may be a 7.2M pixel image capture device with an aspect ratio of about 2:1 (e.g., HxV=3800x1900 pixels) with about 100 degree horizontal FOV. Such an image capture device may be used in place of a three image capture device configuration. Due to significant lens distortion, the vertical FOV of such an image capture device may be significantly less than 50 degrees in implementations in which the image capture device uses a radially symmetric lens. For example, such a lens may not be radially symmetric which would allow for a vertical FOV greater than 50 degrees with 100 degree horizontal FOV.

[0133] The first image capture device 122 may acquire a plurality of first images relative to a scene associated with vehicle 200. Each of the plurality of first images may be acquired as a series of image scan lines, which may be captured using a rolling shutter. Each scan line may include a plurality of pixels.

[0134] The first image capture device 122 may have a scan rate associated with acquisition of each of the first series of image scan lines. The scan rate may refer to a rate at which an image sensor can acquire image data associated with each pixel included in a particular scan line.

[0135] Image capture devices 122, 124, and 126 may contain any suitable type and number of image sensors, including CCD sensors or CMOS sensors, for example. In one embodiment, a CMOS image sensor may be employed along with a rolling shutter, such that each pixel in a row is read one at a time, and scanning of the rows proceeds on a row-by-row basis until an entire image frame has been captured. In some embodiments, the rows may be captured sequentially from top to bottom relative to the frame.

[0136] In some embodiments, one or more of the image capture devices (e.g., image capture devices 122, 124, and 126) disclosed herein may constitute a high resolution imager and may have a resolution greater than 5M pixel, 7M pixel, 10M pixel, or greater.

[0137] The use of a rolling shutter may result in pixels in different rows being exposed and captured at different times, which may cause skew and other image artifacts in the captured image frame. On the other hand, when the image capture device 122 is configured to operate with a global or synchronous shutter, all of the pixels may be exposed for the same amount of time and during a common exposure period. As a result, the image data in a frame collected from a system employing a global shutter represents a snapshot of the entire FOV (such as FOV 202) at a particular time. In contrast, in a rolling shutter application, each row in a frame is exposed and data is capture at different times. Thus, moving objects may appear distorted in an image capture device having a rolling shutter. This phenomenon will be described in greater detail below.

[0138] The second image capture device 124 and the third image capturing device 126 may be any type of image capture device. Like the first image capture device 122, each of image capture devices 124 and 126 may include an optical axis. In one embodiment, each of image capture devices 124 and 126 may include an Aptina M9V024 WVGA sensor with a global shutter. Alternatively, each of image capture devices 124 and 126 may include a rolling shutter. Like image capture device 122, image capture devices 124 and 126 may be configured to include various lenses and optical elements. In some embodiments, lenses associated with image capture devices 124 and 126 may provide FOVs (such as FOVs 204 and 206) that are the same as, or narrower than, a FOV (such as FOV 202) associated with image capture device 122. For example, image capture devices 124 and 126 may have FOVs of 40 degrees, 30 degrees, 26 degrees, 23 degrees, 20 degrees, or less.

[0139] Image capture devices 124 and 126 may acquire a plurality of second and third images relative to a scene associated with vehicle 200. Each of the plurality of second and third images may be acquired as a second and third series of image scan lines, which may be captured using a rolling shutter. Each scan line or row may have a plurality of pixels. Image capture devices 124 and 126 may have second and third scan rates associated with acquisition of each of image scan lines included in the second and third series.

[0140] Each image capture device 122, 124, and 126 may be positioned at any suitable position and orientation relative to vehicle 200. The relative positioning of the image capture devices 122, 124, and 126 may be selected to aid in fusing together the information acquired from the image capture devices. For example, in some embodiments, a FOV (such as FOV 204) associated with image capture device 124 may overlap partially or fully with a FOV (such as FOV 202) associated with image capture device 122 and a FOV (such as FOV 206) associated with image capture device 126.

[0141] Image capture devices 122, 124, and 126 may be located on vehicle 200 at any suitable relative heights. In one instance, there may be a height difference between the image capture devices 122, 124, and 126, which may provide sufficient parallax information to enable stereo analysis. For example, as shown in Fig. 2A, the two image capture devices 122 and 124 are at different heights. There may also be a lateral displacement difference between image capture devices 122, 124, and 126, giving additional parallax information for stereo analysis by processing unit 110, for example. The difference in the lateral displacement may be denoted by d_x , as shown in Figs. 2C and 2D. In some embodiments, fore or aft displacement (e.g., range displacement) may exist between image capture devices 122, 124, and 126. For example, image capture device 122 may be located 0.5 to 2 meters or more behind image capture device 124 and/or image capture device 126. This type of displacement may enable one of the image capture devices to cover potential blind spots of the other image capture device(s).

[0142] Image capture devices 122 may have any suitable resolution capability (e.g., number of pixels associated with the image sensor), and the resolution of the image sensor(s) associated with the image capture device 122 may be higher, lower, or the same as the resolution of the image sensor(s) associated with image capture devices 124 and 126. In some embodiments, the image sensor(s) associated

with image capture device 122 and/or image capture devices 124 and 126 may have a resolution of 640 x 480, 1024 x 768, 1280 x 960, or any other suitable resolution.

[0143] The frame rate (e.g., the rate at which an image capture device acquires a set of pixel data of one image frame before moving on to capture pixel data associated with the next image frame) may be controllable. The frame rate associated with image capture device 122 may be higher, lower, or the same as the frame rate associated with image capture devices 124 and 126. The frame rate associated with image capture devices 122, 124, and 126 may depend on a variety of factors that may affect the timing of the frame rate. For example, one or more of image capture devices 122, 124, and 126 may include a selectable pixel delay period imposed before or after acquisition of image data associated with one or more pixels of an image sensor in image capture device 122, 124, and/or 126. Generally, image data corresponding to each pixel may be acquired according to a clock rate for the device (e.g., one pixel per clock cycle). Additionally, in embodiments including a rolling shutter, one or more of image capture devices 122, 124, and 126 may include a selectable horizontal blanking period imposed before or after acquisition of image data associated with a row of pixels of an image sensor in image capture device 122, 124, and/or 126. Further, one or more of image capture devices 122, 124, and/or 126 may include a selectable vertical blanking period imposed before or after acquisition of image data associated with an image frame of image capture device 122, 124, and 126.

[0144] These timing controls may enable synchronization of frame rates associated with image capture devices 122, 124, and 126, even where the line scan rates of each are different. Additionally, as will be discussed in greater detail below, these selectable timing controls, among other factors (e.g., image sensor resolution, maximum line scan rates, etc.) may enable synchronization of image capture from an area where the FOV of image capture device 122 overlaps with one or more FOVs of image capture devices 124 and 126, even where the field of view of image capture device 122 is different from the FOVs of image capture devices 124 and 126.

[0145] Frame rate timing in image capture device 122, 124, and 126 may depend on the resolution of the associated image sensors. For example, assuming similar line scan rates for both devices, if one device includes an image sensor having a resolution of 640 x 480 and another device includes an image sensor with a resolution of 1280 x 960, then more time will be required to acquire a frame of image data from the sensor having the higher resolution.

[0146] Another factor that may affect the timing of image data acquisition in image capture devices 122, 124, and 126 is the maximum line scan rate. For example, acquisition of a row of image data from an image sensor included in image capture device 122, 124, and 126 will require some minimum amount of time. Assuming no pixel delay periods are added, this minimum amount of time for acquisition of a row of image data will be related to the maximum line scan rate for a particular device. Devices that offer higher maximum line scan rates have the potential to provide higher frame rates than devices with lower maximum line scan rates. In some embodiments, one or more of image capture devices 124 and 126 may have a maximum line scan rate that is higher than a maximum line scan rate associated with image capture device 122. In some embodiments, the maximum line scan rate of image capture device

124 and/or 126 may be 1.25, 1.5, 1.75, or 2 times or more than a maximum line scan rate of image capture device 122.

[0147] In another embodiment, image capture devices 122, 124, and 126 may have the same maximum line scan rate, but image capture device 122 may be operated at a scan rate less than or equal to its maximum scan rate. The system may be configured such that one or more of image capture devices 124 and 126 operate at a line scan rate that is equal to the line scan rate of image capture device 122. In other instances, the system may be configured such that the line scan rate of image capture device 124 and/or image capture device 126 may be 1.25, 1.5, 1.75, or 2 times or more than the line scan rate of image capture device 122.

[0148] In some embodiments, image capture devices 122, 124, and 126 may be asymmetric. That is, they may include cameras having different fields of view (FOV) and focal lengths. The fields of view of image capture devices 122, 124, and 126 may include any desired area relative to an environment of vehicle 200, for example. In some embodiments, one or more of image capture devices 122, 124, and 126 may be configured to acquire image data from an environment in front of vehicle 200, behind vehicle 200, to the sides of vehicle 200, or combinations thereof.

[0149] Further, the focal length associated with each image capture device 122, 124, and/or 126 may be selectable (e.g., by inclusion of appropriate lenses etc.) such that each device acquires images of objects at a desired distance range relative to vehicle 200. For example, in some embodiments image capture devices 122, 124, and 126 may acquire images of close-up objects within a few meters from the vehicle. Image capture devices 122, 124, and 126 may also be configured to acquire images of objects at ranges more distant from the vehicle (e.g., 25 m, 50 m, 100 m, 150 m, or more). Further, the focal lengths of image capture devices 122, 124, and 126 may be selected such that one image capture device (e.g., image capture device 122) can acquire images of objects relatively close to the vehicle (e.g., within 10 m or within 20 m) while the other image capture devices (e.g., image capture devices 124 and 126) can acquire images of more distant objects (e.g., greater than 20 m, 50 m, 100 m, 150 m, etc.) from vehicle 200.

[0150] According to some embodiments, the FOV of one or more image capture devices 122, 124, and 126 may have a wide angle. For example, it may be advantageous to have a FOV of 140 degrees, especially for image capture devices 122, 124, and 126 that may be used to capture images of the area in the vicinity of vehicle 200. For example, image capture device 122 may be used to capture images of the area to the right or left of vehicle 200 and, in such embodiments, it may be desirable for image capture device 122 to have a wide FOV (e.g., at least 140 degrees).

[0151] The field of view associated with each of image capture devices 122, 124, and 126 may depend on the respective focal lengths. For example, as the focal length increases, the corresponding field of view decreases.

[0152] Image capture devices 122, 124, and 126 may be configured to have any suitable fields of view. In one particular example, image capture device 122 may have a horizontal FOV of 46 degrees, image capture device 124 may have a horizontal FOV of 23 degrees, and image capture device 126 may

have a horizontal FOV in between 23 and 46 degrees. In another instance, image capture device 122 may have a horizontal FOV of 52 degrees, image capture device 124 may have a horizontal FOV of 26 degrees, and image capture device 126 may have a horizontal FOV in between 26 and 52 degrees. In some embodiments, a ratio of the FOV of image capture device 122 to the FOVs of image capture device 124 and/or image capture device 126 may vary from 1.5 to 2.0. In other embodiments, this ratio may vary between 1.25 and 2.25.

[0153] System 100 may be configured so that a field of view of image capture device 122 overlaps, at least partially or fully, with a field of view of image capture device 124 and/or image capture device 126. In some embodiments, system 100 may be configured such that the fields of view of image capture devices 124 and 126, for example, fall within (e.g., are narrower than) and share a common center with the field of view of image capture device 122. In other embodiments, the image capture devices 122, 124, and 126 may capture adjacent FOVs or may have partial overlap in their FOVs. In some embodiments, the fields of view of image capture devices 122, 124, and 126 may be aligned such that a center of the narrower FOV image capture devices 124 and/or 126 may be located in a lower half of the field of view of the wider FOV device 122.

[0154] Fig. 2F is a diagrammatic representation of exemplary vehicle control systems, consistent with the disclosed embodiments. As indicated in Fig. 2F, vehicle 200 may include throttling system 220, braking system 230, and steering system 240. System 100 may provide inputs (e.g., control signals) to one or more of throttling system 220, braking system 230, and steering system 240 over one or more data links (e.g., any wired and/or wireless link or links for transmitting data). For example, based on analysis of images acquired by image capture devices 122, 124, and/or 126, system 100 may provide control signals to one or more of throttling system 220, braking system 230, and steering system 240 to navigate vehicle 200 (e.g., by causing an acceleration, a turn, a lane shift, etc.). Further, system 100 may receive inputs from one or more of throttling system 220, braking system 230, and steering system 24 indicating operating conditions of vehicle 200 (e.g., speed, whether vehicle 200 is braking and/or turning, etc.). Further details are provided in connection with Figs. 4-7, below.

[0155] As shown in Fig. 3A, vehicle 200 may also include a user interface 170 for interacting with a driver or a passenger of vehicle 200. For example, user interface 170 in a vehicle application may include a touch screen 320, knobs 330, buttons 340, and a microphone 350. A driver or passenger of vehicle 200 may also use handles (e.g., located on or near the steering column of vehicle 200 including, for example, turn signal handles), buttons (e.g., located on the steering wheel of vehicle 200), and the like, to interact with system 100. In some embodiments, microphone 350 may be positioned adjacent to a rearview mirror 310. Similarly, in some embodiments, image capture device 122 may be located near rearview mirror 310. In some embodiments, user interface 170 may also include one or more speakers 360 (e.g., speakers of a vehicle audio system). For example, system 100 may provide various notifications (e.g., alerts) via speakers 360.

[0156] Figs. 3B-3D are illustrations of an exemplary camera mount 370 configured to be positioned behind a rearview mirror (e.g., rearview mirror 310) and against a vehicle windshield,

consistent with disclosed embodiments. As shown in Fig. 3B, camera mount 370 may include image capture devices 122, 124, and 126. Image capture devices 124 and 126 may be positioned behind a glare shield 380, which may be flush against the vehicle windshield and include a composition of film and/or anti-reflective materials. For example, glare shield 380 may be positioned such that it aligns against a vehicle windshield having a matching slope. In some embodiments, each of image capture devices 122, 124, and 126 may be positioned behind glare shield 380, as depicted, for example, in Fig. 3D. The disclosed embodiments are not limited to any particular configuration of image capture devices 122, 124, and 126, camera mount 370, and glare shield 380. FIG. 3C is an illustration of camera mount 370 shown in Fig. 3B from a front perspective.

[0157] As will be appreciated by a person skilled in the art having the benefit of this disclosure, numerous variations and/or modifications may be made to the foregoing disclosed embodiments. For example, not all components are essential for the operation of system 100. Further, any component may be located in any appropriate part of system 100 and the components may be rearranged into a variety of configurations while providing the functionality of the disclosed embodiments. Therefore, the foregoing configurations are examples and, regardless of the configurations discussed above, system 100 can provide a wide range of functionality to analyze the surroundings of vehicle 200 and navigate vehicle 200 in response to the analysis.

[0158] As discussed below in further detail and consistent with various disclosed embodiments, system 100 may provide a variety of features related to autonomous driving and/or driver assist technology. For example, system 100 may analyze image data, position data (e.g., GPS location information), map data, speed data, and/or data from sensors included in vehicle 200. System 100 may collect the data for analysis from, for example, image acquisition unit 120, position sensor 130, and other sensors. Further, system 100 may analyze the collected data to determine whether or not vehicle 200 should take a certain action, and then automatically take the determined action without human intervention. For example, when vehicle 200 navigates without human intervention, system 100 may automatically control the braking, acceleration, and/or steering of vehicle 200 (e.g., by sending control signals to one or more of throttling system 220, braking system 230, and steering system 240). Further, system 100 may analyze the collected data and issue warnings and/or alerts to vehicle occupants based on the analysis of the collected data. Additional details regarding the various embodiments that are provided by system 100 are provided below.

[0159] Forward-Facing Multi-Imaging System

[0160] As discussed above, system 100 may provide drive assist functionality that uses a multi-camera system. The multi-camera system may use one or more cameras facing in the forward direction of a vehicle. In other embodiments, the multi-camera system may include one or more cameras facing to the side of a vehicle or to the rear of the vehicle. In one embodiment, for example, system 100 may use a two-camera imaging system, where a first camera and a second camera (e.g., image capture devices 122 and 124) may be positioned at the front and/or the sides of a vehicle (e.g., vehicle 200). Other camera configurations are consistent with the disclosed embodiments, and the configurations

disclosed herein are examples. For example, system 100 may include a configuration of any number of cameras (e.g., one, two, three, four, five, six, seven, eight, etc.) Furthermore, system 100 may include “clusters” of cameras. For example, a cluster of cameras (including any appropriate number of cameras, e.g., one, four, eight, etc.) may be forward-facing relative to a vehicle, or may be facing any other
5 direction (e.g., rearward-facing, side-facing, at an angle, etc.) Accordingly, system 100 may include multiple clusters of cameras, with each cluster oriented in a particular direction to capture images from a particular region of a vehicle’s environment.

[0161] The first camera may have a field of view that is greater than, less than, or partially overlapping with, the field of view of the second camera. In addition, the first camera may be connected
10 to a first image processor to perform monocular image analysis of images provided by the first camera, and the second camera may be connected to a second image processor to perform monocular image analysis of images provided by the second camera. The outputs (e.g., processed information) of the first and second image processors may be combined. In some embodiments, the second image processor may receive images from both the first camera and second camera to perform stereo analysis. In another
15 embodiment, system 100 may use a three-camera imaging system where each of the cameras has a different field of view. Such a system may, therefore, make decisions based on information derived from objects located at varying distances both forward and to the sides of the vehicle. References to monocular image analysis may refer to instances where image analysis is performed based on images captured from a single point of view (e.g., from a single camera). Stereo image analysis may refer to instances where
20 image analysis is performed based on two or more images captured with one or more variations of an image capture parameter. For example, captured images suitable for performing stereo image analysis may include images captured: from two or more different positions, from different fields of view, using different focal lengths, along with parallax information, etc.

[0162] For example, in one embodiment, system 100 may implement a three camera
25 configuration using image capture devices 122-126. In such a configuration, image capture device 122 may provide a narrow field of view (e.g., 34 degrees, or other values selected from a range of about 20 to 45 degrees, etc.), image capture device 124 may provide a wide field of view (e.g., 150 degrees or other values selected from a range of about 100 to about 180 degrees), and image capture device 126 may provide an intermediate field of view (e.g., 46 degrees or other values selected from a range of about 35
30 to about 60 degrees). In some embodiments, image capture device 126 may act as a main or primary camera. Image capture devices 122-126 may be positioned behind rearview mirror 310 and positioned substantially side-by-side (e.g., 6 cm apart). Further, in some embodiments, as discussed above, one or more of image capture devices 122-126 may be mounted behind glare shield 380 that is flush with the windshield of vehicle 200. Such shielding may act to minimize the impact of any reflections from inside
35 the car on image capture devices 122-126.

[0163] In another embodiment, as discussed above in connection with Figs. 3B and 3C, the wide field of view camera (e.g., image capture device 124 in the above example) may be mounted lower than the narrow and main field of view cameras (e.g., image devices 122 and 126 in the above example).

This configuration may provide a free line of sight from the wide field of view camera. To reduce reflections, the cameras may be mounted close to the windshield of vehicle 200, and may include polarizers on the cameras to damp reflected light.

[0164] A three camera system may provide certain performance characteristics. For example, some embodiments may include an ability to validate the detection of objects by one camera based on detection results from another camera. In the three camera configuration discussed above, processing unit 110 may include, for example, three processing devices (e.g., three EyeQ series of processor chips, as discussed above), with each processing device dedicated to processing images captured by one or more of image capture devices 122-126.

[0165] In a three camera system, a first processing device may receive images from both the main camera and the narrow field of view camera, and perform vision processing of the narrow FOV camera to, for example, detect other vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. Further, the first processing device may calculate a disparity of pixels between the images from the main camera and the narrow camera and create a 3D reconstruction of the environment of vehicle 200. The first processing device may then combine the 3D reconstruction with 3D map data or with 3D information calculated based on information from another camera.

[0166] The second processing device may receive images from the main camera and perform vision processing to detect other vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. Additionally, the second processing device may calculate a camera displacement and, based on the displacement, calculate a disparity of pixels between successive images and create a 3D reconstruction of the scene (e.g., a structure from motion). The second processing device may send the structure from motion based 3D reconstruction to the first processing device to be combined with the stereo 3D images.

[0167] The third processing device may receive images from the wide FOV camera and process the images to detect vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. The third processing device may further execute additional processing instructions to analyze images to identify objects moving in the image, such as vehicles changing lanes, pedestrians, etc.

[0168] In some embodiments, having streams of image-based information captured and processed independently may provide an opportunity for providing redundancy in the system. Such redundancy may include, for example, using a first image capture device and the images processed from that device to validate and/or supplement information obtained by capturing and processing image information from at least a second image capture device.

[0169] In some embodiments, system 100 may use two image capture devices (e.g., image capture devices 122 and 124) in providing navigation assistance for vehicle 200 and use a third image capture device (e.g., image capture device 126) to provide redundancy and validate the analysis of data received from the other two image capture devices. For example, in such a configuration, image capture devices 122 and 124 may provide images for stereo analysis by system 100 for navigating vehicle 200, while image capture device 126 may provide images for monocular analysis by system 100 to provide

redundancy and validation of information obtained based on images captured from image capture device 122 and/or image capture device 124. That is, image capture device 126 (and a corresponding processing device) may be considered to provide a redundant sub-system for providing a check on the analysis derived from image capture devices 122 and 124 (e.g., to provide an automatic emergency braking (AEB) system). Furthermore, in some embodiments, redundancy and validation of received data may be supplemented based on information received from one more sensors (e.g., radar, lidar, acoustic sensors, information received from one or more transceivers outside of a vehicle, etc.).

[0170] One of skill in the art will recognize that the above camera configurations, camera placements, number of cameras, camera locations, etc., are examples only. These components and others described relative to the overall system may be assembled and used in a variety of different configurations without departing from the scope of the disclosed embodiments. Further details regarding usage of a multi-camera system to provide driver assist and/or autonomous vehicle functionality follow below.

[0171] Fig. 4 is an exemplary functional block diagram of memory 140 and/or 150, which may be stored/programmed with instructions for performing one or more operations consistent with the disclosed embodiments. Although the following refers to memory 140, one of skill in the art will recognize that instructions may be stored in memory 140 and/or 150.

[0172] As shown in Fig. 4, memory 140 may store a monocular image analysis module 402, a stereo image analysis module 404, a velocity and acceleration module 406, and a navigational response module 408. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, applications processor 180 and/or image processor 190 may execute the instructions stored in any of modules 402-408 included in memory 140. One of skill in the art will understand that references in the following discussions to processing unit 110 may refer to applications processor 180 and image processor 190 individually or collectively. Accordingly, steps of any of the following processes may be performed by one or more processing devices.

[0173] In one embodiment, monocular image analysis module 402 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs monocular image analysis of a set of images acquired by one of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from a set of images with additional sensory information (e.g., information from radar) to perform the monocular image analysis. As described in connection with Figs. 5A-5D below, monocular image analysis module 402 may include instructions for detecting a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and any other feature associated with an environment of a vehicle. Based on the analysis, system 100 (e.g., via processing unit 110) may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with navigational response module 408.

[0174] In one embodiment, monocular image analysis module 402 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs monocular

image analysis of a set of images acquired by one of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from a set of images with additional sensory information (e.g., information from radar, lidar, etc.) to perform the monocular image analysis. As described in connection with Figs. 5A-5D below, monocular image analysis module 402 may include

5 instructions for detecting a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and any other feature associated with an environment of a vehicle. Based on the analysis, system 100 (e.g., via processing unit 110) may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with determining a navigational response.

10 [0175] In one embodiment, stereo image analysis module 404 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs stereo image analysis of first and second sets of images acquired by a combination of image capture devices selected from any of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from the first and second sets of images with additional sensory information (e.g.,

15 information from radar) to perform the stereo image analysis. For example, stereo image analysis module 404 may include instructions for performing stereo image analysis based on a first set of images acquired by image capture device 124 and a second set of images acquired by image capture device 126. As described in connection with FIG. 6 below, stereo image analysis module 404 may include instructions for detecting a set of features within the first and second sets of images, such as lane markings, vehicles,

20 pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and the like. Based on the analysis, processing unit 110 may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with navigational response module 408. Furthermore, in some embodiments, stereo image analysis module 404 may implement techniques associated with a trained system (such as a neural network or a deep

25 neural network) or an untrained system.

[0176] In one embodiment, velocity and acceleration module 406 may store software configured to analyze data received from one or more computing and electromechanical devices in vehicle 200 that are configured to cause a change in velocity and/or acceleration of vehicle 200. For example, processing unit 110 may execute instructions associated with velocity and acceleration module

30 406 to calculate a target speed for vehicle 200 based on data derived from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Such data may include, for example, a target position, velocity, and/or acceleration, the position and/or speed of vehicle 200 relative to a nearby vehicle, pedestrian, or road object, position information for vehicle 200 relative to lane markings of the road, and the like. In addition, processing unit 110 may calculate a target speed for vehicle 200 based on

35 sensory input (e.g., information from radar) and input from other systems of vehicle 200, such as throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Based on the calculated target speed, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger a change in velocity and/or

acceleration by, for example, physically depressing the brake or easing up off the accelerator of vehicle 200.

[0177] In one embodiment, navigational response module 408 may store software executable by processing unit 110 to determine a desired navigational response based on data derived from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Such data may include position and speed information associated with nearby vehicles, pedestrians, and road objects, target position information for vehicle 200, and the like. Additionally, in some embodiments, the navigational response may be based (partially or fully) on map data, a predetermined position of vehicle 200, and/or a relative velocity or a relative acceleration between vehicle 200 and one or more objects detected from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Navigational response module 408 may also determine a desired navigational response based on sensory input (e.g., information from radar) and inputs from other systems of vehicle 200, such as throttling system 220, braking system 230, and steering system 240 of vehicle 200. Based on the desired navigational response, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and steering system 240 of vehicle 200 to trigger a desired navigational response by, for example, turning the steering wheel of vehicle 200 to achieve a rotation of a predetermined angle. In some embodiments, processing unit 110 may use the output of navigational response module 408 (e.g., the desired navigational response) as an input to execution of velocity and acceleration module 406 for calculating a change in speed of vehicle 200.

[0178] Furthermore, any of the modules (e.g., modules 402, 404, and 406) disclosed herein may implement techniques associated with a trained system (such as a neural network or a deep neural network) or an untrained system.

[0179] Fig. 5A is a flowchart showing an exemplary process 500A for causing one or more navigational responses based on monocular image analysis, consistent with disclosed embodiments. At step 510, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image acquisition unit 120 (such as image capture device 122 having field of view 202) may capture a plurality of images of an area forward of vehicle 200 (or to the sides or rear of a vehicle, for example) and transmit them over a data connection (e.g., digital, wired, USB, wireless, Bluetooth, etc.) to processing unit 110. Processing unit 110 may execute monocular image analysis module 402 to analyze the plurality of images at step 520, as described in further detail in connection with Figs. 5B-5D below. By performing the analysis, processing unit 110 may detect a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, and the like.

[0180] Processing unit 110 may also execute monocular image analysis module 402 to detect various road hazards at step 520, such as, for example, parts of a truck tire, fallen road signs, loose cargo, small animals, and the like. Road hazards may vary in structure, shape, size, and color, which may make detection of such hazards more challenging. In some embodiments, processing unit 110 may execute monocular image analysis module 402 to perform multi-frame analysis on the plurality of images to

detect road hazards. For example, processing unit 110 may estimate camera motion between consecutive image frames and calculate the disparities in pixels between the frames to construct a 3D-map of the road. Processing unit 110 may then use the 3D-map to detect the road surface, as well as hazards existing above the road surface.

5 [0181] At step 530, processing unit 110 may execute navigational response module 408 to cause one or more navigational responses in vehicle 200 based on the analysis performed at step 520 and the techniques as described above in connection with Fig. 4. Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, and the like. In some embodiments, processing unit 110 may use data derived from execution of velocity and acceleration module 406 to cause the one or
10 more navigational responses. Additionally, multiple navigational responses may occur simultaneously, in sequence, or any combination thereof. For instance, processing unit 110 may cause vehicle 200 to shift one lane over and then accelerate by, for example, sequentially transmitting control signals to steering system 240 and throttling system 220 of vehicle 200. Alternatively, processing unit 110 may cause
15 vehicle 200 to brake while at the same time shifting lanes by, for example, simultaneously transmitting control signals to braking system 230 and steering system 240 of vehicle 200.

 [0182] Fig. 5B is a flowchart showing an exemplary process 500B for detecting one or more vehicles and/or pedestrians in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500B. At step 540, processing unit 110 may determine a set of candidate objects representing possible vehicles and/or
20 pedestrians. For example, processing unit 110 may scan one or more images, compare the images to one or more predetermined patterns, and identify within each image possible locations that may contain objects of interest (e.g., vehicles, pedestrians, or portions thereof). The predetermined patterns may be designed in such a way to achieve a high rate of “false hits” and a low rate of “misses.” For example, processing unit 110 may use a low threshold of similarity to predetermined patterns for identifying
25 candidate objects as possible vehicles or pedestrians. Doing so may allow processing unit 110 to reduce the probability of missing (e.g., not identifying) a candidate object representing a vehicle or pedestrian.

 [0183] At step 542, processing unit 110 may filter the set of candidate objects to exclude certain candidates (e.g., irrelevant or less relevant objects) based on classification criteria. Such criteria may be derived from various properties associated with object types stored in a database (e.g., a database
30 stored in memory 140). Properties may include object shape, dimensions, texture, position (e.g., relative to vehicle 200), and the like. Thus, processing unit 110 may use one or more sets of criteria to reject false candidates from the set of candidate objects.

 [0184] At step 544, processing unit 110 may analyze multiple frames of images to determine whether objects in the set of candidate objects represent vehicles and/or pedestrians. For example,
35 processing unit 110 may track a detected candidate object across consecutive frames and accumulate frame-by-frame data associated with the detected object (e.g., size, position relative to vehicle 200, etc.). Additionally, processing unit 110 may estimate parameters for the detected object and compare the object’s frame-by-frame position data to a predicted position.

[0185] At step 546, processing unit 110 may construct a set of measurements for the detected objects. Such measurements may include, for example, position, velocity, and acceleration values (relative to vehicle 200) associated with the detected objects. In some embodiments, processing unit 110 may construct the measurements based on estimation techniques using a series of time-based observations such as Kalman filters or linear quadratic estimation (LQE), and/or based on available modeling data for different object types (e.g., cars, trucks, pedestrians, bicycles, road signs, etc.). The Kalman filters may be based on a measurement of an object's scale, where the scale measurement is proportional to a time to collision (e.g., the amount of time for vehicle 200 to reach the object). Thus, by performing steps 540-546, processing unit 110 may identify vehicles and pedestrians appearing within the set of captured images and derive information (e.g., position, speed, size) associated with the vehicles and pedestrians. Based on the identification and the derived information, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A, above.

[0186] At step 548, processing unit 110 may perform an optical flow analysis of one or more images to reduce the probabilities of detecting a "false hit" and missing a candidate object that represents a vehicle or pedestrian. The optical flow analysis may refer to, for example, analyzing motion patterns relative to vehicle 200 in the one or more images associated with other vehicles and pedestrians, and that are distinct from road surface motion. Processing unit 110 may calculate the motion of candidate objects by observing the different positions of the objects across multiple image frames, which are captured at different times. Processing unit 110 may use the position and time values as inputs into mathematical models for calculating the motion of the candidate objects. Thus, optical flow analysis may provide another method of detecting vehicles and pedestrians that are nearby vehicle 200. Processing unit 110 may perform optical flow analysis in combination with steps 540-546 to provide redundancy for detecting vehicles and pedestrians and increase the reliability of system 100.

[0187] Fig. 5C is a flowchart showing an exemplary process 500C for detecting road marks and/or lane geometry information in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500C. At step 550, processing unit 110 may detect a set of objects by scanning one or more images. To detect segments of lane markings, lane geometry information, and other pertinent road marks, processing unit 110 may filter the set of objects to exclude those determined to be irrelevant (e.g., minor potholes, small rocks, etc.). At step 552, processing unit 110 may group together the segments detected in step 550 belonging to the same road mark or lane mark. Based on the grouping, processing unit 110 may develop a model to represent the detected segments, such as a mathematical model.

[0188] At step 554, processing unit 110 may construct a set of measurements associated with the detected segments. In some embodiments, processing unit 110 may create a projection of the detected segments from the image plane onto the real-world plane. The projection may be characterized using a 3rd-degree polynomial having coefficients corresponding to physical properties such as the position, slope, curvature, and curvature derivative of the detected road. In generating the projection, processing unit 110 may take into account changes in the road surface, as well as pitch and roll rates associated with

vehicle 200. In addition, processing unit 110 may model the road elevation by analyzing position and motion cues present on the road surface. Further, processing unit 110 may estimate the pitch and roll rates associated with vehicle 200 by tracking a set of feature points in the one or more images.

[0189] At step 556, processing unit 110 may perform multi-frame analysis by, for example, tracking the detected segments across consecutive image frames and accumulating frame-by-frame data associated with detected segments. As processing unit 110 performs multi-frame analysis, the set of measurements constructed at step 554 may become more reliable and associated with an increasingly higher confidence level. Thus, by performing steps 550-556, processing unit 110 may identify road marks appearing within the set of captured images and derive lane geometry information. Based on the identification and the derived information, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with Fig. 5A, above.

[0190] At step 558, processing unit 110 may consider additional sources of information to further develop a safety model for vehicle 200 in the context of its surroundings. Processing unit 110 may use the safety model to define a context in which system 100 may execute autonomous control of vehicle 200 in a safe manner. To develop the safety model, in some embodiments, processing unit 110 may consider the position and motion of other vehicles, the detected road edges and barriers, and/or general road shape descriptions extracted from map data (such as data from map database 160). By considering additional sources of information, processing unit 110 may provide redundancy for detecting road marks and lane geometry and increase the reliability of system 100.

[0191] Fig. 5D is a flowchart showing an exemplary process 500D for detecting traffic lights in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500D. At step 560, processing unit 110 may scan the set of images and identify objects appearing at locations in the images likely to contain traffic lights. For example, processing unit 110 may filter the identified objects to construct a set of candidate objects, excluding those objects unlikely to correspond to traffic lights. The filtering may be done based on various properties associated with traffic lights, such as shape, dimensions, texture, position (e.g., relative to vehicle 200), and the like. Such properties may be based on multiple examples of traffic lights and traffic control signals and stored in a database. In some embodiments, processing unit 110 may perform multi-frame analysis on the set of candidate objects reflecting possible traffic lights. For example, processing unit 110 may track the candidate objects across consecutive image frames, estimate the real-world position of the candidate objects, and filter out those objects that are moving (which are unlikely to be traffic lights). In some embodiments, processing unit 110 may perform color analysis on the candidate objects and identify the relative position of the detected colors appearing inside possible traffic lights.

[0192] At step 562, processing unit 110 may analyze the geometry of a junction. The analysis may be based on any combination of: (i) the number of lanes detected on either side of vehicle 200, (ii) markings (such as arrow marks) detected on the road, and (iii) descriptions of the junction extracted from map data (such as data from map database 160). Processing unit 110 may conduct the analysis using information derived from execution of monocular analysis module 402. In addition, Processing unit 110

may determine a correspondence between the traffic lights detected at step 560 and the lanes appearing near vehicle 200.

[0193] As vehicle 200 approaches the junction, at step 564, processing unit 110 may update the confidence level associated with the analyzed junction geometry and the detected traffic lights. For instance, the number of traffic lights estimated to appear at the junction as compared with the number actually appearing at the junction may impact the confidence level. Thus, based on the confidence level, processing unit 110 may delegate control to the driver of vehicle 200 in order to improve safety conditions. By performing steps 560-564, processing unit 110 may identify traffic lights appearing within the set of captured images and analyze junction geometry information. Based on the identification and the analysis, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A, above.

[0194] Fig. 5E is a flowchart showing an exemplary process 500E for causing one or more navigational responses in vehicle 200 based on a vehicle path, consistent with the disclosed embodiments. At step 570, processing unit 110 may construct an initial vehicle path associated with vehicle 200. The vehicle path may be represented using a set of points expressed in coordinates (x, z) , and the distance d_i between two points in the set of points may fall in the range of 1 to 5 meters. In one embodiment, processing unit 110 may construct the initial vehicle path using two polynomials, such as left and right road polynomials. Processing unit 110 may calculate the geometric midpoint between the two polynomials and offset each point included in the resultant vehicle path by a predetermined offset (e.g., a smart lane offset), if any (an offset of zero may correspond to travel in the middle of a lane). The offset may be in a direction perpendicular to a segment between any two points in the vehicle path. In another embodiment, processing unit 110 may use one polynomial and an estimated lane width to offset each point of the vehicle path by half the estimated lane width plus a predetermined offset (e.g., a smart lane offset).

[0195] At step 572, processing unit 110 may update the vehicle path constructed at step 570. Processing unit 110 may reconstruct the vehicle path constructed at step 570 using a higher resolution, such that the distance d_k between two points in the set of points representing the vehicle path is less than the distance d_i described above. For example, the distance d_k may fall in the range of 0.1 to 0.3 meters. Processing unit 110 may reconstruct the vehicle path using a parabolic spline algorithm, which may yield a cumulative distance vector S corresponding to the total length of the vehicle path (i.e., based on the set of points representing the vehicle path).

[0196] At step 574, processing unit 110 may determine a look-ahead point (expressed in coordinates as (x_l, z_l)) based on the updated vehicle path constructed at step 572. Processing unit 110 may extract the look-ahead point from the cumulative distance vector S , and the look-ahead point may be associated with a look-ahead distance and look-ahead time. The look-ahead distance, which may have a lower bound ranging from 10 to 20 meters, may be calculated as the product of the speed of vehicle 200 and the look-ahead time. For example, as the speed of vehicle 200 decreases, the look-ahead distance may also decrease (e.g., until it reaches the lower bound). The look-ahead time, which may range from 0.5 to

1.5 seconds, may be inversely proportional to the gain of one or more control loops associated with causing a navigational response in vehicle 200, such as the heading error tracking control loop. For example, the gain of the heading error tracking control loop may depend on the bandwidth of a yaw rate loop, a steering actuator loop, car lateral dynamics, and the like. Thus, the higher the gain of the heading error tracking control loop, the lower the look-ahead time.

[0197] At step 576, processing unit 110 may determine a heading error and yaw rate command based on the look-ahead point determined at step 574. Processing unit 110 may determine the heading error by calculating the arctangent of the look-ahead point, e.g., $\arctan(x_l/z_l)$. Processing unit 110 may determine the yaw rate command as the product of the heading error and a high-level control gain. The high-level control gain may be equal to: $(2 / \text{look-ahead time})$, if the look-ahead distance is not at the lower bound. Otherwise, the high-level control gain may be equal to: $(2 * \text{speed of vehicle 200} / \text{look-ahead distance})$.

[0198] Fig. 5F is a flowchart showing an exemplary process 500F for determining whether a leading vehicle is changing lanes, consistent with the disclosed embodiments. At step 580, processing unit 110 may determine navigation information associated with a leading vehicle (e.g., a vehicle traveling ahead of vehicle 200). For example, processing unit 110 may determine the position, velocity (e.g., direction and speed), and/or acceleration of the leading vehicle, using the techniques described in connection with FIGS. 5A and 5B, above. Processing unit 110 may also determine one or more road polynomials, a look-ahead point (associated with vehicle 200), and/or a snail trail (e.g., a set of points describing a path taken by the leading vehicle), using the techniques described in connection with FIG. 5E, above.

[0199] At step 582, processing unit 110 may analyze the navigation information determined at step 580. In one embodiment, processing unit 110 may calculate the distance between a snail trail and a road polynomial (e.g., along the trail). If the variance of this distance along the trail exceeds a predetermined threshold (for example, 0.1 to 0.2 meters on a straight road, 0.3 to 0.4 meters on a moderately curvy road, and 0.5 to 0.6 meters on a road with sharp curves), processing unit 110 may determine that the leading vehicle is likely changing lanes. In the case where multiple vehicles are detected traveling ahead of vehicle 200, processing unit 110 may compare the snail trails associated with each vehicle. Based on the comparison, processing unit 110 may determine that a vehicle whose snail trail does not match with the snail trails of the other vehicles is likely changing lanes. Processing unit 110 may additionally compare the curvature of the snail trail (associated with the leading vehicle) with the expected curvature of the road segment in which the leading vehicle is traveling. The expected curvature may be extracted from map data (e.g., data from map database 160), from road polynomials, from other vehicles' snail trails, from prior knowledge about the road, and the like. If the difference in curvature of the snail trail and the expected curvature of the road segment exceeds a predetermined threshold, processing unit 110 may determine that the leading vehicle is likely changing lanes.

[0200] In another embodiment, processing unit 110 may compare the leading vehicle's instantaneous position with the look-ahead point (associated with vehicle 200) over a specific period of

time (e.g., 0.5 to 1.5 seconds). If the distance between the leading vehicle's instantaneous position and the look-ahead point varies during the specific period of time, and the cumulative sum of variation exceeds a predetermined threshold (for example, 0.3 to 0.4 meters on a straight road, 0.7 to 0.8 meters on a moderately curvy road, and 1.3 to 1.7 meters on a road with sharp curves), processing unit 110 may
5 determine that the leading vehicle is likely changing lanes. In another embodiment, processing unit 110 may analyze the geometry of the snail trail by comparing the lateral distance traveled along the trail with the expected curvature of the snail trail. The expected radius of curvature may be determined according to the calculation: $(\delta_z^2 + \delta_x^2) / 2 / (\delta_x)$, where δ_x represents the lateral distance traveled and δ_z represents the longitudinal distance traveled. If the difference between the lateral distance traveled and the expected
10 curvature exceeds a predetermined threshold (e.g., 500 to 700 meters), processing unit 110 may determine that the leading vehicle is likely changing lanes. In another embodiment, processing unit 110 may analyze the position of the leading vehicle. If the position of the leading vehicle obscures a road polynomial (e.g., the leading vehicle is overlaid on top of the road polynomial), then processing unit 110 may determine that the leading vehicle is likely changing lanes. In the case where the position of the leading vehicle is
15 such that, another vehicle is detected ahead of the leading vehicle and the snail trails of the two vehicles are not parallel, processing unit 110 may determine that the (closer) leading vehicle is likely changing lanes.

[0201] At step 584, processing unit 110 may determine whether or not leading vehicle 200 is changing lanes based on the analysis performed at step 582. For example, processing unit 110 may make
20 the determination based on a weighted average of the individual analyses performed at step 582. Under such a scheme, for example, a decision by processing unit 110 that the leading vehicle is likely changing lanes based on a particular type of analysis may be assigned a value of "1" (and "0" to represent a determination that the leading vehicle is not likely changing lanes). Different analyses performed at step 582 may be assigned different weights, and the disclosed embodiments are not limited to any particular
25 combination of analyses and weights. Furthermore, in some embodiments, the analysis may make use of trained system (e.g., a machine learning or deep learning system), which may, for example, estimate a future path ahead of a current location of a vehicle based on an image captured at the current location.

[0202] Fig. 6 is a flowchart showing an exemplary process 600 for causing one or more navigational responses based on stereo image analysis, consistent with disclosed embodiments. At step
30 610, processing unit 110 may receive a first and second plurality of images via data interface 128. For example, cameras included in image acquisition unit 120 (such as image capture devices 122 and 124 having fields of view 202 and 204) may capture a first and second plurality of images of an area forward of vehicle 200 and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive the first and second plurality
35 of images via two or more data interfaces. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0203] At step 620, processing unit 110 may execute stereo image analysis module 404 to perform stereo image analysis of the first and second plurality of images to create a 3D map of the road in

front of the vehicle and detect features within the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, and the like. Stereo image analysis may be performed in a manner similar to the steps described in connection with FIGS. 5A-5D, above. For example, processing unit 110 may execute stereo image analysis module 404 to detect candidate objects (e.g., vehicles, pedestrians, road marks, traffic lights, road hazards, etc.) within the first and second plurality of images, filter out a subset of the candidate objects based on various criteria, and perform multi-frame analysis, construct measurements, and determine a confidence level for the remaining candidate objects. In performing the steps above, processing unit 110 may consider information from both the first and second plurality of images, rather than information from one set of images alone. For example, processing unit 110 may analyze the differences in pixel-level data (or other data subsets from among the two streams of captured images) for a candidate object appearing in both the first and second plurality of images. As another example, processing unit 110 may estimate a position and/or velocity of a candidate object (e.g., relative to vehicle 200) by observing that the object appears in one of the plurality of images but not the other or relative to other differences that may exist relative to objects appearing in the two image streams. For example, position, velocity, and/or acceleration relative to vehicle 200 may be determined based on trajectories, positions, movement characteristics, etc. of features associated with an object appearing in one or both of the image streams.

[0204] At step 630, processing unit 110 may execute navigational response module 408 to cause one or more navigational responses in vehicle 200 based on the analysis performed at step 620 and the techniques as described above in connection with FIG. 4. Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, a change in velocity, braking, and the like. In some embodiments, processing unit 110 may use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. Additionally, multiple navigational responses may occur simultaneously, in sequence, or any combination thereof.

[0205] Fig. 7 is a flowchart showing an exemplary process 700 for causing one or more navigational responses based on an analysis of three sets of images, consistent with disclosed embodiments. At step 710, processing unit 110 may receive a first, second, and third plurality of images via data interface 128. For instance, cameras included in image acquisition unit 120 (such as image capture devices 122, 124, and 126 having fields of view 202, 204, and 206) may capture a first, second, and third plurality of images of an area forward and/or to the side of vehicle 200 and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive the first, second, and third plurality of images via three or more data interfaces. For example, each of image capture devices 122, 124, 126 may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0206] At step 720, processing unit 110 may analyze the first, second, and third plurality of images to detect features within the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, and the like. The analysis may be performed in a manner

similar to the steps described in connection with Figs. 5A-5D and 6, above. For instance, processing unit 110 may perform monocular image analysis (e.g., via execution of monocular image analysis module 402 and based on the steps described in connection with Figs. 5A-5D, above) on each of the first, second, and third plurality of images. Alternatively, processing unit 110 may perform stereo image analysis (e.g., via execution of stereo image analysis module 404 and based on the steps described in connection with FIG. 6, above) on the first and second plurality of images, the second and third plurality of images, and/or the first and third plurality of images. The processed information corresponding to the analysis of the first, second, and/or third plurality of images may be combined. In some embodiments, processing unit 110 may perform a combination of monocular and stereo image analyses. For example, processing unit 110 may perform monocular image analysis (e.g., via execution of monocular image analysis module 402) on the first plurality of images and stereo image analysis (e.g., via execution of stereo image analysis module 404) on the second and third plurality of images. The configuration of image capture devices 122, 124, and 126—including their respective locations and fields of view 202, 204, and 206—may influence the types of analyses conducted on the first, second, and third plurality of images. The disclosed embodiments are not limited to a particular configuration of image capture devices 122, 124, and 126, or the types of analyses conducted on the first, second, and third plurality of images.

[0207] In some embodiments, processing unit 110 may perform testing on system 100 based on the images acquired and analyzed at steps 710 and 720. Such testing may provide an indicator of the overall performance of system 100 for certain configurations of image capture devices 122, 124, and 126. For example, processing unit 110 may determine the proportion of “false hits” (e.g., cases where system 100 incorrectly determined the presence of a vehicle or pedestrian) and “misses.”

[0208] At step 730, processing unit 110 may cause one or more navigational responses in vehicle 200 based on information derived from two of the first, second, and third plurality of images. Selection of two of the first, second, and third plurality of images may depend on various factors, such as, for example, the number, types, and sizes of objects detected in each of the plurality of images. Processing unit 110 may also make the selection based on image quality and resolution, the effective field of view reflected in the images, the number of captured frames, the extent to which one or more objects of interest actually appear in the frames (e.g., the percentage of frames in which an object appears, the proportion of the object that appears in each such frame, etc.), and the like.

[0209] In some embodiments, processing unit 110 may select information derived from two of the first, second, and third plurality of images by determining the extent to which information derived from one image source is consistent with information derived from other image sources. For example, processing unit 110 may combine the processed information derived from each of image capture devices 122, 124, and 126 (whether by monocular analysis, stereo analysis, or any combination of the two) and determine visual indicators (e.g., lane markings, a detected vehicle and its location and/or path, a detected traffic light, etc.) that are consistent across the images captured from each of image capture devices 122, 124, and 126. Processing unit 110 may also exclude information that is inconsistent across the captured images (e.g., a vehicle changing lanes, a lane model indicating a vehicle that is too close to vehicle 200,

etc.). Thus, processing unit 110 may select information derived from two of the first, second, and third plurality of images based on the determinations of consistent and inconsistent information.

[0210] Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, and the like. Processing unit 110 may cause the one or more navigational responses based on the analysis performed at step 720 and the techniques as described above in connection with FIG. 4. Processing unit 110 may also use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. In some embodiments, processing unit 110 may cause the one or more navigational responses based on a relative position, relative velocity, and/or relative acceleration between vehicle 200 and an object detected within any of the first, second, and third plurality of images. Multiple navigational responses may occur simultaneously, in sequence, or any combination thereof.

[0211] Reinforcement Learning and Trained Navigational Systems

[0212] The sections that follow discuss autonomous driving along with systems and methods for accomplishing autonomous control of a vehicle, whether that control is fully autonomous (a self-driving vehicle) or partially autonomous (e.g., one or more driver assist systems or functions). As shown in Fig. 8, the autonomous driving task can be partitioned into three main modules, including a sensing module 801, a driving policy module 803, and a control module 805. In some embodiments, modules 801, 803, and 805 may be stored in memory unit 140 and/or memory unit 150 of system 100, or modules 801, 803, and 805 (or portions thereof) may be stored remotely from system 100 (e.g., stored in a server accessible to system 100 via, for example, wireless transceiver 172). Furthermore, any of the modules (e.g., modules 801, 803, and 805) disclosed herein may implement techniques associated with a trained system (such as a neural network or a deep neural network) or an untrained system.

[0213] Sensing module 801, which may be implemented using processing unit 110, may handle various tasks relating to sensing of a navigational state in an environment of a host vehicle. Such tasks may rely upon input from various sensors and sensing systems associated with the host vehicle. These inputs may include images or image streams from one or more onboard cameras, GPS position information, accelerometer outputs, user feedback, or user inputs to one or more user interface devices, radar, lidar, etc. Sensing, which may include data from cameras and/or any other available sensors, along with map information, may be collected, analyzed, and formulated into a “sensed state,” describing information extracted from a scene in the environment of the host vehicle. The sensed state may include sensed information relating to target vehicles, lane markings, pedestrians, traffic lights, road geometry, lane shape, obstacles, distances to other objects/vehicles, relative velocities, relative accelerations, among any other potential sensed information. Supervised machine learning may be implemented in order to produce a sensing state output based on sensed data provided to sensing module 801. The output of the sensing module may represent a sensed navigational “state” of the host vehicle, which may be passed to driving policy module 803.

[0214] While a sensed state may be developed based on image data received from one or more cameras or image sensors associated with a host vehicle, a sensed state for use in navigation may be

developed using any suitable sensor or combination of sensors. In some embodiments, the sensed state may be developed without reliance upon captured image data. In fact, any of the navigational principles described herein may be applicable to sensed states developed based on captured image data as well as sensed states developed using other non-image based sensors. The sensed state may also be determined via sources external to the host vehicle. For example, a sensed state may be developed in full or in part based on information received from sources remote from the host vehicle (e.g., based on sensor information, processed state information, etc. shared from other vehicles, shared from a central server, or from any other source of information relevant to a navigational state of the host vehicle.)

[0215] Driving policy module 803, which is discussed in more detail below and which may be implemented using processing unit 110, may implement a desired driving policy in order to decide on one or more navigational actions for the host vehicle to take in response to the sensed navigational state. If there are no other agents (e.g., target vehicles or pedestrians) present in the environment of the host vehicle, the sensed state input to driving policy module 803 may be handled in a relatively straightforward manner. The task becomes more complex when the sensed state requires negotiation with one or more other agents. The technology used to generate the output of driving policy module 803 may include reinforcement learning (discussed in more detail below). The output of driving policy module 803 may include at least one navigational action for the host vehicle and may include a desired acceleration (which may translate to an updated speed for the host vehicle), a desired yaw rate for the host vehicle, a desired trajectory, among other potential desired navigational actions.

[0216] Based on the output from the driving policy module 803, control module 805, which may also be implemented using processing unit 110, may develop control instructions for one or more actuators or controlled devices associated with the host vehicle. Such actuators and devices may include an accelerator, one or more steering controls, a brake, a signal transmitter, a display, or any other actuator or device that may be controlled as part of a navigation operation associated with a host vehicle. Aspects of control theory may be used to generate the output of control module 805. Control module 805 may be responsible for developing and outputting instructions to controllable components of the host vehicle in order to implement the desired navigational goals or requirements of driving policy module 803.

[0217] Returning to driving policy module 803, in some embodiments, a trained system trained through reinforcement learning may be used to implement driving policy module 803. In other embodiments, driving policy module 803 may be implemented without a machine learning approach, by using specified algorithms to “manually” address the various scenarios that may arise during autonomous navigation. Such an approach, however, while viable, may result in a driving policy that is too simplistic and may lack the flexibility of a trained system based on machine learning. A trained system, for example, may be better equipped to handle complex navigational states and may better determine whether a taxi is parking or is stopping to pick up or drop off a passenger; determine whether a pedestrian intends to cross the street ahead of the host vehicle; balance unexpected behavior of other drivers with defensiveness; negotiate in dense traffic involving target vehicles and/or pedestrians; decide when to suspend certain navigational rules or augment other rules; anticipate unsensed, but anticipated conditions

(e.g., whether a pedestrian will emerge from behind a car or obstacle); etc. A trained system based on reinforcement learning may also be better equipped to address a state space that is continuous and high-dimensional along with an action space that is continuous.

[0218] Training of the system using reinforcement learning may involve learning a driving policy in order to map from sensed states to navigational actions. A driving policy is a function
 5 $\pi : S \rightarrow A$, where S is a set of states and $A \subset \mathbb{R}^2$ is the action space (e.g., desired speed, acceleration, yaw commands, etc.). The state space is $S = S_s \times S_p$, where S_s is the sensing state and S_p is additional information on the state saved by the policy. Working in discrete time intervals, at time t , the current state $s_t \in S$ may be observed, and the policy may be applied to obtain a desired action,
 10 $a_t = \pi(s_t)$.

[0219] The system may be trained through exposure to various navigational states, having the system apply the policy, providing a reward (based on a reward function designed to reward desirable navigational behavior). Based on the reward feedback, the system may “learn” the policy and becomes trained in producing desirable navigational actions. For example, the learning system may observe the
 15 current state $s_t \in S$ and decide on an action $a_t \in A$ based on a policy $\pi : S \rightarrow \mathbb{D}(A)$. Based on the decided action (and implementation of the action), the environment moves to the next state $s_{t+1} \in S$ for observation by the learning system. For each action developed in response to the observed state, the feedback to the learning system is a reward signal r_1, r_2, \dots .

[0220] The goal of Reinforcement Learning (RL) is to find a policy π . It is usually assumed
 20 that at time t , there is a reward function r_t which measures the instantaneous quality of being at state s_t and taking action a_t . However, taking the action a_t at time t affects the environment and therefore affects the value of the future states. As a result, when deciding on what action to take, not only should the current reward be taken into account, but future rewards should also be considered. In some instances the system should take a certain action, even though it is associated with a reward lower than another
 25 available option, when the system determines that in the future a greater reward may be realized if the lower reward option is taken now. To formalize this, observe that a policy, π , and an initial state, s , induces a distribution over \mathbb{R}^T , where the probability of a vector (r_1, \dots, r_T) is the probability of observing the rewards r_1, \dots, r_T , if the agent starts at state $s_0 = s$ and from there on follows the policy π . The value of the initial state s may be defined as:

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=1}^T r_t \mid s_0 = s, \forall t \geq 1, a_t = \pi(s_t) \right].$$

30 [0221]

[0222] Instead of restricting the time horizon to T , the future rewards may be discounted to define, for some fixed $\gamma \in (0, 1)$:

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=1}^{\infty} \gamma^t r_t \mid s_0 = s, \forall t \geq 1, a_t = \pi(s_t) \right].$$

[0223]

[0224] In any case, the optimal policy is the solution of

$$\operatorname{argmax}_{\pi} \mathbb{E}[V^{\pi}(s)]$$

[0225]

[0226] where the expectation is over the initial state, s .

[0227] There are several possible methodologies for training the driving policy system. For
 5 example, an imitation approach (e.g., behavior cloning) may be used in which the system learns from
 state/action pairs where the actions are those that would be chosen by a good agent (e.g., a human) in
 response to a particular observed state. Suppose a human driver is observed. Through this observation,
 many examples of the form (s_t, a_t) , where s_t is the state and a_t is the action of the human driver could be
 obtained, observed, and used as a basis for training the driving policy system. For example, supervised
 10 learning can be used to learn a policy π such that $\pi(s_t) \approx a_t$. There are many potential advantages of
 this approach. First, there is no requirement to define a reward function. Second, the learning is
 supervised and happens offline (there is no need to apply the agent in the learning process). A
 disadvantage of this method is that different human drivers, and even the same human drivers, are not
 deterministic in their policy choices. Hence, learning a function for which $\|\pi(s_t) - a_t\|$ is very small
 15 is often infeasible. And, even small errors may accumulate over time to yield large errors.

[0228] Another technique that may be employed is policy based learning. Here, the policy
 may be expressed in parametric form and directly optimized using a suitable optimization technique (e.g.,

stochastic gradient descent). The approach is to directly solve the problem given in $\operatorname{argmax}_{\pi} \mathbb{E}[V^{\pi}(s)]$.

There are of course many ways to solve the problem. One advantage of this approach is that it tackles the
 20 problem directly, and therefore often leads to good practical results. One potential disadvantage is that it
 often requires an “on-policy” training, namely, the learning of π is an iterative process, where at iteration j
 we have a non-perfect policy, π_j , and to construct the next policy π_j , we must interact with the
 environment while acting based on π_j .

[0229] The system may also be trained through value based learning (learning Q or V
 25 functions). Suppose a good approximation can be learned to the optimal value function V^* . An optimal
 policy may be constructed (e.g., by relying on the Bellman equation). Some versions of value based
 learning can be implemented offline (called “off-policy” training). Some disadvantages of the value-based
 approach may result from its strong dependence on Markovian assumptions and required approximation
 of a complicated function (it may be more difficult to approximate the value function than to approximate
 30 the policy directly).

[0230] Another technique may include model based learning and planning (learning the
 probability of state transitions and solving the optimization problem of finding the optimal V).
 Combinations of these techniques may also be used to train the learning system. In this approach, the
 dynamics of the process may be learned, namely, the function that takes (s_t, a_t) and yields a distribution
 35 over the next state s_{t+1} . Once this function is learned, the optimization problem may be solved to find the
 policy π whose value is optimal. This is called “planning”. One advantage of this approach may be that

the learning part is supervised and can be applied offline by observing triplets (s_t, a_t, s_{t+1}) . One disadvantage of this approach, similar to the “imitation” approach, may be that small errors in the learning process can accumulate and to yield inadequately performing policies.

[0231] Another approach for training driving policy module 803 may include decomposing the driving policy function into semantically meaningful components. This allows implementation of parts of the policy manually, which may ensure the safety of the policy, and implementation of other parts of the policy using reinforcement learning techniques, which may enable adaptivity to many scenarios, a human-like balance between defensive/aggressive behavior, and a human-like negotiation with other drivers. From the technical perspective, a reinforcement learning approach may combine several methodologies and offer a tractable training procedure, where most of the training can be performed using either recorded data or a self-constructed simulator.

[0232] In some embodiments, training of driving policy module 803 may rely upon an “options” mechanism. To illustrate, consider a simple scenario of a driving policy for a two-lane highway. In a direct RL approach, a policy π that maps the state into $A \subset \mathbb{R}^2$, where the first component of $\pi(s)$ is the desired acceleration command and the second component of $\pi(s)$ is the yaw rate. In a modified approach, the following policies can be constructed:

[0233] Automatic Cruise Control (ACC) policy, $o_{ACC} : S \rightarrow A$: this policy always outputs a yaw rate of 0 and only changes the speed so as to implement smooth and accident-free driving.

[0234] ACC+Left policy, $o_L : S \rightarrow A$: the longitudinal command of this policy is the same as the ACC command. The yaw rate is a straightforward implementation of centering the vehicle toward the middle of the left lane, while ensuring a safe lateral movement (e.g., don’t move left if there’s a car on the left side).

[0235] ACC+Right policy, $o_R : S \rightarrow A$: Same as o_L , but the vehicle may be centered toward the middle of the right lane.

[0236] These policies may be referred to as “options”. Relying on these “options”, a policy can be learned that selects options, $\pi_o : S \rightarrow O$, where O is the set of available options. In one case, $O = \{o_{ACC}, o_L, o_R\}$. The option-selector policy, π_o , defines an actual policy, $\pi : S \rightarrow A$, by setting, for every s , $\pi(s) = o_{\pi_o(s)}(s)$.

[0237] In practice, the policy function may be decomposed into an options graph 901, as shown in Fig. 9. Another example options graph 1000 is shown in Fig. 10. The options graph can represent a hierarchical set of decisions organized as a Directed Acyclic Graph (DAG). There is a special node called the root node 903 of the graph. This node has no incoming nodes. The decision process traverses through the graph, starting from the root node, until it reaches a “leaf” node, which refers to a node that has no outgoing decision lines. As shown in Fig. 9, leaf nodes may include nodes 905, 907, and 909, for example. Upon encountering a leaf node, driving policy module 803 may output the acceleration and steering commands associated with a desired navigational action associated with the leaf node.

[0238] Internal nodes, such as nodes 911, 913, and 915, for example, may result in implementation of a policy that chooses a child among its available options. The set of available children of an internal node include all of the nodes associated with a particular internal node via decision lines. For example, internal node 913 designated as “Merge” in Fig. 9 includes three children nodes 909, 915, and 917 (“Stay,” “Overtake Right,” and “Overtake Left,” respectively) each joined to node 913 by a decision line.

[0239] Flexibility of the decision-making system may be gained by enabling nodes to adjust their position in the hierarchy of the options graph. For example, any of the nodes may be allowed to declare themselves as “critical.” Each node may implement a function “is critical,” that outputs “True” if the node is in a critical section of its policy implementation. For example, a node that is responsible for a take-over, may declare itself as critical while in the middle of a maneuver. This may impose constraints on the set of available children of a node u , which may include all nodes v which are children of node u and for which there exists a path from v to a leaf node that goes through all nodes designated as critical. Such an approach may allow, on one hand, declaration of the desired path on the graph at each time step, while on the other hand, stability of a policy may be preserved, especially while critical portions of the policy are being implemented.

[0240] By defining an options graph, the problem of learning the driving policy $\pi : S \rightarrow A$ may be decomposed into a problem of defining a policy for each node of the graph, where the policy at internal nodes should choose from among available children nodes. For some of the nodes, the respective policy may be implemented manually (e.g., through if-then type algorithms specifying a set of actions in response to an observed state) while for others the policies may be implemented using a trained system built through reinforcement learning. The choice between manual or trained/learned approaches may depend on safety aspects associated with the task and on its relative simplicity. The option graphs may be constructed in a manner such that some of the nodes are straightforward to implement, while other nodes may rely on trained models. Such an approach can ensure safe operation of the system.

[0241] The following discussion provides further details regarding the role of the options graph of Fig. 9 within driving policy module 803. As discussed above, the input to the driving policy module is a “sensed state,” which summarizes the environment map, for example, as obtained from available sensors. The output of driving policy module 803 is a set of desires (optionally, together with a set of hard constraints) that define a trajectory as a solution of an optimization problem.

[0242] As described above, the options graph represents a hierarchical set of decisions organized as a DAG. There is a special node called the “root” of the graph. The root node is the only node that has no incoming edges (e.g., decision lines). The decision process traverses the graph, starting from the root node, until it reaches a “leaf” node, namely, a node that has no outgoing edges. Each internal node should implement a policy that picks a child among its available children. Every leaf node should implement a policy that, based on the entire path from the root to the leaf, defines a set of Desires (e.g., a set of navigational goals for the host vehicle). The set of Desires, together with a set of hard constraints that are defined directly based on the sensed state, establish an optimization problem whose solution is the

trajectory for the vehicle. The hard constraints may be employed to further increase the safety of the system, and the Desires can be used to provide driving comfort and human-like driving behavior of the system. The trajectory provided as a solution to the optimization problem, in turn, defines the commands that should be provided to the steering, braking, and/or engine actuators in order to accomplish the trajectory.

[0243] Returning to Fig. 9, options graph 901 represents an options graph for a two-lane highway, including with merging lanes (meaning that at some points, a third lane is merged into either the right or the left lane of the highway). The root node 903 first decides if the host vehicle is in a plain road scenario or approaching a merge scenario. This is an example of a decision that can be implemented based on the sensing state. Plain road node 911 includes three child nodes: stay node 909, overtake left node 917, and overtake right node 915. Stay refers to a situation in which the host vehicle would like to keep driving in the same lane. The stay node is a leaf node (no outgoing edges/lines). Therefore, it the stay node defines a set of Desires. The first Desire it defines may include the desired lateral position — e.g., as close as possible to the center of the current lane of travel. There may also be a desire to navigate smoothly (e.g., within predefined or allowable acceleration maximums). The stay node may also define how the host vehicle is to react to other vehicles. For example, the stay node may review sensed target vehicles and assign each a semantic meaning, which can be translated into components of the trajectory.

[0244] Various semantic meanings may be assigned to target vehicles in an environment of the host vehicle. For example, in some embodiments the semantic meaning may include any of the following designations: 1) not relevant: indicating that the sensed vehicle in the scene is currently not relevant; 2) next lane: indicating that the sensed vehicle is in an adjacent lane and an appropriate offset should be maintained relative to this vehicle (the exact offset may be calculated in the optimization problem that constructs the trajectory given the Desires and hard constraints, and can potentially be vehicle dependent—the stay leaf of the options graph sets the target vehicle's semantic type, which defines the Desire relative to the target vehicle); 3) give way: the host vehicle will attempt to give way to the sensed target vehicle by, for example, reducing speed (especially where the host vehicle determines that the target vehicle is likely to cut into the lane of the host vehicle); 4) take way: the host vehicle will attempt to take the right of way by, for example, increasing speed; 5) follow: the host vehicle desires to maintain smooth driving following after this target vehicle; 6) takeover left/right: this means the host vehicle would like to initiate a lane change to the left or right lane. Overtake left node 917 and overtake right node 915 are internal nodes that do not yet define Desires.

[0245] The next node in options graph 901 is the select gap node 919. This node may be responsible for selecting a gap between two target vehicles in a particular target lane that host vehicle desires to enter. By choosing a node of the form ID_j, for some value of j, the host vehicle arrives at a leaf that designates a Desire for the trajectory optimization problem—e.g., the host vehicle wishes to make a maneuver so as to arrive at the selected gap. Such a maneuver may involve first accelerating/braking in the current lane and then heading to the target lane at an appropriate time to enter the selected gap. If the

select gap node 919 cannot find an appropriate gap, it moves to the abort node 921, which defines a desire to move back to the center of the current lane and cancel the takeover.

[0246] Returning to merge node 913, when the host vehicle approaches a merge, it has several options that may depend on a particular situation. For example, as shown in Fig. 11A, host vehicle 1105 is traveling along a two-lane road with no other target vehicles detected, either in the primary lanes of the two-lane road or in the merge lane 1111. In this situation, driving policy module 803, upon reaching merge node 913, may select stay node 909. That is, staying within its current lane may be desired where no target vehicles are sensed as merging onto the roadway.

[0247] In Fig. 11B, the situation is slightly different. Here, host vehicle 1105 senses one or more target vehicles 1107 entering the main roadway 1112 from merge lane 1111. In this situation, once driving policy module 803 encounters merge node 913, it may choose to initiate an overtake left maneuver in order to avoid the merging situation.

[0248] In Fig. 11C, host vehicle 1105 encounters one or more target vehicles 1107 entering main roadway 1112 from merge lane 1111. Host vehicle 1105 also detects target vehicles 1109 traveling in a lane adjacent to the lane of the host vehicle. The host vehicle also detects one or more target vehicles 1110 traveling in the same lane as host vehicle 1105. In this situation, driving policy module 803 may decide to adjust the speed of host vehicle 1105 to give way to target vehicle 1107 and to proceed ahead of target vehicle 1115. This can be accomplished, for example, by progressing to select gap node 919, which, in turn, will select a gap between ID0 (vehicle 1107) and ID1 (vehicle 1115) as the appropriate merging gap. In such a case, the appropriate gap of the merging situation defines the objective for a trajectory planner optimization problem.

[0249] As discussed above, nodes of the options graph may declare themselves as “critical,” which may ensure that the selected option passes through the critical nodes. Formally, each node may implement a function IsCritical. After performing a forward pass on the options graph, from the root to a leaf, and solving the optimization problem of the trajectory planner, a backward pass may be performed from the leaf back to the root. Along this backward pass, the IsCritical function of all nodes in the pass may be called, and a list of all critical nodes may be saved. In the forward path corresponding to the next time frame, driving policy module 803 may be required to choose a path from the root node to a leaf that goes through all critical nodes.

[0250] Figs. 11A-11C may be used to show a potential benefit of this approach. For example, in a situation where an overtake action is initiated, and driving policy module 803 arrives at the leaf corresponding to ID_k, it would be undesirable to choose, for example, the stay node 909 when the host vehicle is in the middle of the takeover maneuver. To avoid such jumpiness, the ID_j node can designate itself as critical. During the maneuver, the success of the trajectory planner can be monitored, and function IsCritical will return a “True” value if the overtake maneuver progresses as intended. This approach may ensure that in the next time frame, the takeover maneuver will be continued (rather than jumping to another, potentially inconsistent maneuver prior to completion of the initially selected maneuver). If, on the other hand, monitoring of the maneuver indicates that the selected maneuver is not

progressing as intended, or if the maneuver has become unnecessary or impossible, the function `IsCritical` can return a “False” value. This can allow the select gap node to select a different gap in the next time frame, or to abort the overtake maneuver altogether. This approach may allow, on one hand, declaration of the desired path on the options graph at each time step, while on the other hand, may help to promote stability of the policy while in critical parts of the execution.

[0251] Hard constraints, which will be discussed in more detail below, may be differentiated from navigational desires. For example, hard constraints may ensure safe driving by applying an added layer of filtering of a planned navigational action. The implicated hard constraints, which may be programmed and defined manually, rather than through use of a trained system built upon reinforcement learning, can be determined from the sensed state. In some embodiments, however, the trained system may learn the applicable hard constraints to be applied and followed. Such an approach may promote driving policy module 803 arriving at a selected action that is already in compliance with the applicable hard constraints, which may reduce or eliminate selected actions that may require later modification to comply with applicable hard constraints. Nevertheless, as a redundant safety measure, hard constraints may be applied to the output of driving policy module 803 even where driving policy module 803 has been trained to account for predetermined hard constraints.

[0252] There are many examples of potential hard constraints. For example, a hard constraint may be defined in conjunction with a guardrail on an edge of a road. In no situation may the host vehicle be allowed to pass the guardrail. Such a rule induces a hard lateral constraint on the trajectory of the host vehicle. Another example of a hard constraint may include a road bump (e.g., a speed control bump), which may induce a hard constraint on the speed of driving before the bump and while traversing the bump. Hard constraints may be considered safety critical and, therefore, may be defined manually rather than relying solely on a trained system learning the constraints during training.

[0253] In contrast to hard constraints, the goal of desires may be to enable or achieve comfortable driving. As discussed above, an example of a desire may include a goal of positioning the host vehicle at a lateral position within a lane that corresponds to the center of the host vehicle lane. Another desire may include the ID of a gap to fit into. Note that there is not a requirement for the host vehicle to be exactly in the center of the lane, but instead a desire to be as close as possible to it may ensure that the host vehicle tends to migrate to the center of the lane even in the event of deviations from the center of the lane. Desires may not be safety critical. In some embodiments, desires may require negotiation with other drivers and pedestrians. One approach for constructing the desires may rely on the options graph, and the policy implemented in at least some nodes of the graph may be based on reinforcement learning.

[0254] For the nodes of options graph 901 or 1000 implemented as nodes trained based on learning, the training process may include decomposing the problem into a supervised learning phase and a reinforcement learning phase. In the supervised learning phase, a differentiable mapping from (s_t, a_t) to \hat{s}_{t+1} can be learned such that $\hat{s}_{t+1} \approx s_{t+1}$. This may be similar to “model-based” reinforcement learning. However, in the forward loop of the network, \hat{s}_{t+1} may be replaced by the actual value of s_{t+1} , therefore

eliminating the problem of error accumulation. The role of prediction of \hat{s}_{t+1} is to propagate messages from the future back to past actions. In this sense, the algorithm may be a combination of “model-based” reinforcement learning with “policy-based learning.”

[0255] An important element that may be provided in some scenarios is a differentiable path
5 from future losses/rewards back to decisions on actions. With the option graph structure, the implementation of options that involve safety constraints are usually not differentiable. To overcome this issue, the choice of a child in a learned policy node may be stochastic. That is, a node may output a probability vector, p , that assigns probabilities used in choosing each of the children of the particular node. Suppose that a node has k children and let $a^{(1)}, \dots, a^{(k)}$ be the actions of the path from each child
10 to a leaf. The resulting predicted action is therefore $\hat{a} = \sum_{i=1}^k p_i a^{(i)}$, which may result in a differentiable path from the action to p . In practice, an action a may be chosen to be $a^{(i)}$ for $i \sim p$, and the difference between a and \hat{a} may be referred to as additive noise.

[0256] For the training of \hat{s}_{t+1} given s_t, a_t , supervised learning may be used together with real data. For training the policy of nodes simulators can be used. Later, fine tuning of a policy can be
15 accomplished using real data. Two concepts may make the simulation more realistic. First, using imitation, an initial policy can be constructed using the “behavior cloning” paradigm, using large real-world data sets. In some cases, the resulting agents may be suitable. In other cases, the resulting agents at least form very good initial policies for the other agents on the roads. Second, using self-play, our own policy may be used to augment the training. For example, given an initial implementation of the other
20 agents (cars/pedestrians) that may be experienced, a policy may be trained based on a simulator. Some of the other agents may be replaced with the new policy, and the process may be repeated. As a result, the policy can continue to improve as it should respond to a larger variety of other agents that have differing levels of sophistication.

[0257] Further, in some embodiments, the system may implement a multi-agent approach. For
25 example, the system may take into account data from various sources and/or images capturing from multiple angles. Further, some disclosed embodiments may provide economy of energy, as anticipation of an event which does not directly involve the host vehicle, but which may have an effect on the host vehicle can be considered, or even anticipation of an event that may lead to unpredictable circumstances involving other vehicles may be a consideration (e.g., radar may “see through” the leading vehicle and
30 anticipation of an unavoidable, or even a high likelihood of an event that will affect the host vehicle).

[0258] Trained System with Imposed Navigational Constraints

[0259] In the context of autonomous driving, a significant concern is how to ensure that a learned policy of a trained navigational network will be safe. In some embodiments, the driving policy system may be trained using constraints, such that the actions selected by the trained system may already
35 account for applicable safety constraints. Additionally, in some embodiments, an extra layer of safety may be provided by passing the selected actions of the trained system through one or more hard constraints implicated by a particular sensed scene in the environment of the host vehicle. Such an

approach may ensure that that the actions taken by the host vehicle have been restricted to those confirmed as satisfying applicable safety constraints.

[0260] At its core, the navigational system may include a learning algorithm based on a policy function that maps an observed state to one or more desired actions. In some implementations, the learning algorithm is a deep learning algorithm. The desired actions may include at least one action expected to maximize an anticipated reward for a vehicle. While in some cases, the actual action taken by the vehicle may correspond to one of the desired actions, in other cases, the actual action taken may be determined based on the observed state, one or more desired actions, and non-learned, hard constraints (e.g., safety constraints) imposed on the learning navigational engine. These constraints may include no-drive zones surrounding various types of detected objects (e.g., target vehicles, pedestrians, stationary objects on the side of a road or in a roadway, moving objects on the side of a road or in a roadway, guard rails, etc.) In some cases, the size of the zone may vary based on a detected motion (e.g., speed and/or direction) of a detected object. Other constraints may include a maximum speed of travel when passing within an influence zone of a pedestrian, a maximum deceleration (to account for a target vehicle spacing behind the host vehicle), a mandatory stop at a sensed crosswalk or railroad crossing, etc.

[0261] Hard constraints used in conjunction with a system trained through machine learning may offer a degree of safety in autonomous driving that may surpass a degree of safety available based on the output of the trained system alone. For example, the machine learning system may be trained using a desired set of constraints as training guidelines and, therefore, the trained system may select an action in response to a sensed navigational state that accounts for and adheres to the limitations of applicable navigational constraints. Still, however, the trained system has some flexibility in selecting navigational actions and, therefore, there may exist at least some situations in which an action selected by the trained system may not strictly adhere to relevant navigational constraints. Therefore, in order to require that a selected action strictly adheres to relevant navigational constraints, the output of the trained system may be combined with, compared to, filtered with, adjusted, modified, etc. using a non-machine learning component outside the learning/trained framework that guarantees strict application of relevant navigational constraints.

[0262] The following discussion provides additional details regarding the trained system and the potential benefits (especially from a safety perspective) that may be gleaned from combining a trained system with an algorithmic component outside of the trained/learning framework. As discussed, the reinforcement learning objective by policy may be optimized through stochastic gradient ascent. The objective (e.g., the expected reward) may be defined as $\mathbb{E}_{\bar{s} \sim P_n} R(\bar{s})$.

[0263] Objectives that involve expectation may be used in machine learning scenarios. Such an objective, without being bound by navigational constraints, however, may not return actions strictly bound by those constraints. For example, considering a reward function for which $R(\bar{s}) = -r$ for trajectories that represent a rare "corner" event to be avoided (e.g., such as an accident), and $R(\bar{s}) \in [-1, 1]$ for the rest of the trajectories, one goal for the learning system may be to learn to perform an overtake maneuver. Normally, in an accident free trajectory, $R(\bar{s})$ would reward successful, smooth,

takeovers and penalize staying in a lane without completing the takeover- hence the range $[-1, 1]$. If a sequence, \bar{s} , represents an accident, the reward, $-r$, should provide a sufficiently high penalty to discourage such an occurrence. The question is what should be the value of r to ensure accident-free driving.

5 [0264] Observe that the effect of an accident on $\mathbb{E}[R(\bar{s})]$ is the additive term $-pr$ where p is the probability mass of trajectories with an accident event. If this term is negligible, i.e., $p \ll 1/r$, then the learning system may prefer a policy that performs an accident (or adopt in general a reckless driving policy) in order to fulfill the takeover maneuver successfully more often than a policy that would be more defensive at the expense of having some takeover maneuvers not complete successfully. In other words, if the probability of accidents is to be at most p , then r must be set such that $r \gg 1/p$. It may be desirable to make p extremely small (e.g., on the order of $p = 10^{-9}$). Therefore, r should be large. In policy gradient, the gradient of $\mathbb{E}[R(\bar{s})]$ may be estimated. The following lemma shows that the variance of the random variable $R(\bar{s})$ grows with pr^2 , which is larger than r for $r \gg 1/p$. Therefore, estimating the objective may be difficult, and estimating its gradient may be even more difficult.

15 [0265] Lemma: Let π_0 be a policy and let p and r be scalars such that with probability p , $R(\bar{s}) = -r$ is obtained, and with probability $1-p$ we have $R(\bar{s}) \in [-1, 1]$ is obtained. Then,

$$[0266] \quad \text{Var}[R(\bar{s})] \geq pr^2 - (pr + (1-p))^2 = (p-p^2)r^2 - 2p(1-p)r - (1-p)^2 \approx pr^2$$

[0267] where the last approximation holds for the case $r \geq 1/p$.

20 [0268] This discussion shows that an objection of the form $\mathbb{E}[R(\bar{s})]$ may not ensure functional safety without causing a variance problem. The baseline subtraction method for variance reduction may not offer a sufficient remedy to the problem because the problem would shift from a high variance of $R(\bar{s})$ to an equally high variance of the baseline constants whose estimation would equally suffer numeric instabilities. Moreover, if the probability of an accident is p , then on average at least $1/p$ sequences should be sampled before obtaining an accident event. This implies a lower bound of $1/p$ samples of sequences for a learning algorithm that aims at minimizing $\mathbb{E}[R(\bar{s})]$. The solution to this problem may be found in the architectural design described herein, rather than through numerical conditioning techniques. The approach here is based on the notion that hard constraints should be injected outside of the learning framework. In other words, the policy function may be decomposed into a learnable part and a non-learnable part. Formally, the policy function may be structured as

30 $\pi_\theta = \pi^{(T)} \circ \pi_\theta^{(D)}$, where $\pi_\theta^{(D)}$ maps the (agnostic) state space into a set of Desires (e.g., desired navigational goals, etc.), while $\pi^{(T)}$ maps the Desires into a trajectory (which may determine how the car should move in a short range). The function $\pi_\theta^{(D)}$ is responsible for the comfort of driving and for making strategic decisions such as which other cars should be over-taken or given way and what is the

desired position of the host vehicle within its lane, etc. The mapping from sensed navigational state to Desires is a policy $\pi_{\theta}^{(D)}$ that may be learned from experience by maximizing an expected reward. The desires produced by $\pi_{\theta}^{(D)}$ may be translated into a cost function over driving trajectories. The function $\pi^{(T)}$, not a learned function, may be implemented by finding a trajectory that minimizes the cost subject to hard constraints on functional safety. This decomposition may ensure functional safety while at the same time providing for comfortable driving.

[0269] A double merge navigational situation, as depicted in Fig. 11D, provides an example further illustrating these concepts. In a double merge, vehicles approach the merge area 1130 from both left and right sides. And, from each side, a vehicle, such as vehicle 1133 or vehicle 1135, can decide whether to merge into lanes on the other side of merge area 1130. Successfully executing a double merge in busy traffic may require significant negotiation skills and experience and may be difficult to execute in a heuristic or brute force approach by enumerating all possible trajectories that could be taken by all agents in the scene. In this double merge example, a set of Desires, \mathcal{D} , appropriate for the double merge maneuver may be defined. \mathcal{D} may be the Cartesian product of the following sets:

$\mathcal{D} = [0, v_{\max}] \times L \times \{g, t, o\}^n$, where $[0, v_{\max}]$ is the desired target speed of the host vehicle, $L = \{1, 1.5, 2, 2.5, 3, 3.5, 4\}$ is the desired lateral position in lane units where whole numbers designate a lane center and fractional numbers designate lane boundaries, and $\{g, t, o\}$ are classification labels assigned to each of the n other vehicles. The other vehicles may be assigned “g” if the host vehicle is to give way to the other vehicle, “t” if the host vehicle is to take way relative to the other vehicle, or “o” if the host vehicle is to maintain an offset distance relative to the other vehicle.

[0270] Below is a description of how a set of Desires, $(v, l, c_1, \dots, c_n) \in \mathcal{D}$, may be translated into a cost function over driving trajectories. A driving trajectory may be represented by $(x_1, y_1), \dots, (x_k, y_k)$, where (x_i, y_i) is the (lateral, longitudinal) location of the host vehicle (in ego-centric units) at time $\tau \cdot i$. In some experiments, $\tau = 0.1\text{sec}$ and $k = 10$. Of course, other values may be selected as well. The cost assigned to a trajectory may include a weighted sum of individual costs assigned to the desired speed, lateral position, and the label assigned to each of the other n vehicles.

[0271] Given a desired speed $v \in [0, v_{\max}]$, the cost of a trajectory associated with speed is

$$\sum_{i=2}^k (v - \|(x_i, y_i) - (x_{i-1}, y_{i-1})\|/\tau)^2.$$

[0272] Given desired lateral position, $l \in L$, the cost associated with desired lateral position is

$$\sum_{i=1}^k \text{dist}(x_i, y_i, l).$$

[0273] where $\text{dist}(x, y, l)$ is the distance from the point (x, y) to the lane position l . Regarding the cost due to other vehicles, for any other vehicle $(x'_1, y'_1), \dots, (x'_k, y'_k)$ may represent the other vehicle in egocentric units of the host vehicle, and i may be the earliest point for which there exists j such that the distance between (x_i, y_i) and (x'_j, y'_j) is small. If there is no such point, then i can be set as $i = \infty$.

5 If another car is classified as "give-way", it may be desirable that $\tau i > \tau j + 0.5$, meaning that the host vehicle will arrive to the trajectory intersection point at least 0.5 seconds after the other vehicle will arrive at the same point. A possible formula for translating the above constraint into a cost is $[\tau(j - i) + 0.5]_+$.

[0274] Likewise, if another car is classified as "take-way", it may be desirable that $\tau j > \tau i + 0.5$, which may be translated to the cost $[\tau(i - j) + 0.5]_+$. If another car is classified as "offset", it may be desirable that $i = \infty$, meaning that the trajectory of the host vehicle and the trajectory of the offset car do not intersect. This condition can be translated to a cost by penalizing with respect to the distance between trajectories.

[0275] Assigning a weight to each of these costs may provide a single objective function for the trajectory planner, $\pi^{(T)}$. A cost that encourages smooth driving may be added to the objective. And, to ensure functional safety of the trajectory, hard constraints can be added to the objective. For example, (x_i, y_i) may be prohibited from being off the roadway, and (x_i, y_i) may be forbidden from being close to (x'_j, y'_j) for any trajectory point (x'_j, y'_j) of any other vehicle if $|i - j|$ is small.

[0276] To summarize, the policy, π_θ , can be decomposed into a mapping from the agnostic state to a set of Desires and a mapping from the Desires to an actual trajectory. The latter mapping is not based on learning and may be implemented by solving an optimization problem whose cost depends on the Desires and whose hard constraints may guarantee functional safety of the policy.

[0277] The following discussion describes mapping from the agnostic state to the set of Desires. As described above, to be compliant with functional safety, a system reliant upon reinforcement learning alone may suffer a high and unwieldy variance on the reward $R(\bar{s})$. This result may be avoided by decomposing the problem into a mapping from (agnostic) state space to a set of Desires using policy gradient iterations followed by a mapping to an actual trajectory which does not involve a system trained based on machine learning.

[0278] For various reasons, the decision making may be further decomposed into semantically meaningful components. For example, the size of \mathcal{D} might be large and even continuous. In the double-merge scenario described above with respect to Fig. 11D, $\mathcal{D} = [0, v_{\max}] \times L \times \{g, t, o\}^n$.

Additionally, the gradient estimator may involve the term $\sum_{t=1}^T \nabla_{\theta} \pi_{\theta}(a_t | s_t)$. In such an expression, the variance may grow with the time horizon T . In some cases, the value of T may be roughly 250 which may be high enough to create significant variance. Supposing a sampling rate is in the range of 10Hz and the merge area 1130 is 100 meters, preparation for the merge may begin

approximately 300 meters before the merge area. If the host vehicle travels at 16 meters per second (about 60 km per hour), then the value of T for an episode may be roughly 250.

[0279] Returning to the concept of an options graph, an options graph that may be representative of the double merge scenario depicted in Fig. 11D is shown in Fig. 11E. As previously discussed, an options graph may represent a hierarchical set of decisions organized as a Directed Acyclic Graph (DAG). There may be a special node in the graph called the "root" node 1140, which may be the only node that has no incoming edges (e.g., decision lines). The decision process may traverse the graph, starting from the root node, until it reaches a "leaf" node, namely, a node that has no outgoing edges. Each internal node may implement a policy function that chooses a child from among its available children. There may be a predefined mapping from the set of traversals over the options graph to the set of desires, \mathcal{D} . In other words, a traversal on the options graph may be automatically translated into a desire in \mathcal{D} . Given a node, v , in the graph, a parameter vector θ_v may specify the policy of choosing a child of v . If θ is the concatenation of all the θ_v , then $\pi_{\theta}^{(D)}$ may be defined by traversing from the root of the graph to a leaf, while at each node v using the policy defined by θ_v , to choose a child node.

[0280] In the double merge options graph 1139 of Fig. 11E, root node 1140 may first decide if the host vehicle is within the merging area (e.g., area 1130 of Fig. 11D) or if the host vehicle instead is approaching the merging area and needs to prepare for a possible merge. In both cases, the host vehicle may need to decide whether to change lanes (e.g., to the left or to the right side) or whether to stay in the current lane. If the host vehicle has decided to change lanes, the host vehicle may need to decide whether conditions are suitable to go on and perform the lane change maneuver (e.g., at "go" node 1142). If it is not possible to change lanes, the host vehicle may attempt to "push" toward the desired lane (e.g., at node 1144 as part of a negotiation with vehicles in the desired lane) by aiming at being on the lane mark. Alternatively, the host vehicle may opt to "stay" in the same lane (e.g., at node 1146). Such a process may determine the lateral position for the host vehicle in a natural way. For example,

[0281] This may enable determination of the desired lateral position in a natural way. For example, if the host vehicle changes lanes from lane 2 to lane 3, the "go" node may set the desired lateral position to 3, the "stay" node may set the desired lateral position to 2, and the "push" node may set the desired lateral position to 2.5. Next, the host vehicle may decide whether to maintain the "same" speed (node 1148), "accelerate" (node 1150), or "decelerate" (node 1152). Next, the host vehicle may enter a "chain like" structure 1154 that goes over the other vehicles and sets their semantic meaning to a value in the set $\{g, t, o\}$. This process may set the desires relative to the other vehicles. The parameters of all nodes in this chain may be shared (similar to Recurrent Neural Networks).

[0282] A potential benefit of the options is the interpretability of the results. Another potential benefit is that the decomposable structure of the set \mathcal{D} can be relied upon and, therefore, the policy at each node may be chosen from among a small number of possibilities. Additionally, the structure may allow for a reduction in the variance of the policy gradient estimator.

[0283] As discussed above, the length of an episode in the double merge scenario may be roughly $T = 250$ steps. Such a value (or any other suitable value depending on a particular navigational scenario) may provide enough time to see the consequences of the host vehicle actions (e.g., if the host vehicle decided to change lanes as a preparation for the merge, the host vehicle will see the benefit only after a successful completion of the merge). On the other hand, due to the dynamic of driving, the host vehicle must make decisions at a fast enough frequency (e.g., 10Hz in the case described above).

[0284] The options graph may enable a decrease in the effective value of T in at least two ways. First, given higher level decisions, a reward can be defined for lower level decisions while taking into account shorter episodes. For example, when the host vehicle has already chosen a "lane change" and the "go" node, a policy can be learned for assigning semantic meaning to vehicles by looking at episodes of 2-3 seconds (meaning that T becomes 20 - 30 instead of 250). Second, for high level decisions (such as whether to change lanes or to stay in the same lane), the host vehicle may not need to make decisions every 0.1 seconds. Instead, the host vehicle may be able to either make decisions at a lower frequency (e.g., every second), or implement an "option termination" function, and then the gradient may be calculated only after every termination of the option. In both cases, the effective value of T may be an order of magnitude smaller than its original value. All in all, the estimator at every node may depend on a value of T which is an order of magnitude smaller than the original 250 steps, which may immediately transfer to a smaller variance.

[0285] As discussed above, hard constraints may promote safer driving, and there may be several different types of constraints. For example, static hard constraints may be defined directly from the sensing state. These may include speed bumps, speed limits, road curvature, junctions, etc., within the environment of the host vehicle that may implicate one or more constraints on vehicle speed, heading, acceleration, braking (deceleration), etc. Static hard constraints may also include semantic free space where the host vehicle is prohibited from going outside of the free space and from navigating too close to physical barriers, for example. Static hard constraints may also limit (e.g., prohibit) maneuvers that do not comply with various aspects of a kinematic motion of the vehicle, for example, a static hard constraint can be used to prohibit maneuvers that might lead to the host vehicle overturning, sliding, or otherwise losing control.

[0286] Hard constraints may also be associated with vehicles. For example, a constraint may be employed requiring that a vehicle maintain a longitudinal distance to other vehicles of at least one meter and a lateral distance from other vehicles of at least 0.5 meters. Constraints may also be applied such that the host vehicle will avoid maintaining a collision course with one or more other vehicles. For example, a time τ may be a measure of time based on a particular scene. The predicted trajectories of the host vehicle and one or more other vehicles may be considered from a current time to time τ . Where the two trajectories intersect, (t_i^a, t_i^l) may represent the time of arrival and the leaving time of vehicle i to the intersection point. That is, each car will arrive at point when a first part of the car passes the intersection point, and a certain amount of time will be required before the last part of the car passes through the intersection point. This amount of time separates the arrival time from the leaving time. Assuming that,

$t_1^a < t_2^a$ (i.e., that the arrival time of vehicle 1 is less than the arrival time of vehicle 2), then we will want to ensure that vehicle 1 has left the intersection point prior to vehicle 2 arriving. Otherwise, a collision would result. Thus, a hard constraint may be implemented such that $t_1^l > t_2^a$. Moreover, to ensure that vehicle 1 and vehicle 2 do not miss one another by a minimal amount, an added margin of safety may be obtained by including a buffer time into the constraint (e.g., 0.5 seconds or another suitable value). A hard constraint relating to predicted intersection trajectories of two vehicles may be expressed as $t_1^l > t_2^a + 0.5$.

[0287] The amount of time τ over which the trajectories of the host vehicle and one or more other vehicles are tracked may vary. In junction scenarios, however, where speeds may be lower, τ may be longer, and τ may be defined such that a host vehicle will enter and leave the junction in less than τ seconds.

[0288] Applying hard constraints to vehicle trajectories, of course, requires that the trajectories of those vehicles be predicted. For the host vehicle, trajectory prediction may be relatively straightforward, as the host vehicle generally already understands and, indeed, is planning an intended trajectory at any given time. Relative to other vehicles, predicting their trajectories can be less straightforward. For other vehicles, the baseline calculation for determining predicted trajectories may rely on the current speed and heading of the other vehicles, as determined, for example, based on analysis of an image stream captured by one or more cameras and/or other sensors (radar, lidar, acoustic, etc.) aboard the host vehicle.

[0289] There can be some exceptions, however, that can simplify the problem or at least provide added confidence in a trajectory predicted for another vehicle. For example, with respect to structured roads in which there is an indication of lanes and where give-way rules may exist, the trajectories of other vehicles can be based, at least in part, upon the position of the other vehicles relative to the lanes and based upon applicable give-way rules. Thus, in some situations, when there are observed lane structures, it may be assumed that next-lane vehicles will respect lane boundaries. That is, the host vehicle may assume that a next-lane vehicle will stay in its lane unless there is observed evidence (e.g., a signal light, strong lateral movement, movement across a lane boundary) indicating that the next-lane vehicle will cut into the lane of the host vehicle.

[0290] Other situations may also provide clues regarding the expected trajectories of other vehicles. For example, at stop signs, traffic lights, roundabouts, etc., where the host vehicle may have the right of way, it may be assumed that other vehicles will respect that right of way. Thus, unless there is observed evidence of a rule break, other vehicles may be assumed to proceed along a trajectory that respects the rights of way possessed by the host vehicle.

[0291] Hard constraints may also be applied with respect to pedestrians in an environment of the host vehicle. For example, a buffer distance may be established with respect to pedestrians such that the host vehicle is prohibited from navigating any closer than the prescribed buffer distance relative to any observed pedestrian. The pedestrian buffer distance may be any suitable distance. In some embodiments, the buffer distance may be at least one meter relative to an observed pedestrian.

[0292] Similar to the situation with vehicles, hard constraints may also be applied with respect to relative motion between pedestrians and the host vehicle. For example, the trajectory of a pedestrian (based on a heading direction and speed) may be monitored relative to the projected trajectory of the host vehicle. Given a particular pedestrian trajectory, with every point p on the trajectory, $t(p)$ may represent the time required for the pedestrian to reach point p . To maintain the required buffer distance of at least 1 meter from the pedestrian, either $t(p)$ must be larger than the time the host vehicle will reach point p (with sufficient difference in time such that the host vehicle passes in front of the pedestrian by a distance of at least one meter) or that $t(p)$ must be less than the time the host vehicle will reach point p (e.g., if the host vehicle brakes to give way to the pedestrian). Still, in the latter example, the hard constraint may require that the host vehicle arrive at point p at a sufficient time later than the pedestrian such that the host vehicle can pass behind the pedestrian and maintain the required buffer distance of at least one meter. Of course, there may be exceptions to the pedestrian hard constraint. For example, where the host vehicle has the right of way or where speeds are very slow, and there is no observed evidence that the pedestrian will decline to give way to the host vehicle or will otherwise navigate toward the host vehicle, the pedestrian hard constraint may be relaxed (e.g., to a smaller buffer of at least 0.75 meters or 0.50 meters).

[0293] In some examples, constraints may be relaxed where it is determined that not all can be met. For example, in situations where a road is too narrow to leave desired spacing (e.g., 0.5 meters) from both curbs or from a curb and a parked vehicle, one or more the constraints may be relaxed if there are mitigating circumstances. For example, if there are no pedestrians (or other objects) on the sidewalk one can proceed slowly at 0.1 meters from a curb. In some embodiments, constraints may be relaxed if doing so will improve the user experience. For example, in order to avoid a pothole, constraints may be relaxed to allow a vehicle to navigate closer to the edges of the lane, a curb, or a pedestrian more than might ordinarily be permitted. Furthermore, when determining which constraints to relax, in some embodiments, the one or more constraints chosen to relax are those deemed to have the least available negative impact to safety. For example, a constraint relating to how close the vehicle may travel to the curb or to a concrete barrier may be relaxed before relaxing one dealing with proximity to other vehicles. In some embodiments, pedestrian constraints may be the last to be relaxed, or may never be relaxed in some situations.

[0294] Fig. 12 shows an example of a scene that may be captured and analyzed during navigation of a host vehicle. For example, a host vehicle may include a navigation system (e.g., system 100), as described above, that may receive from a camera (e.g., at least one of image capture device 122, image capture device 124, and image capture device 126) associated with the host vehicle a plurality of images representative of an environment of the host vehicle. The scene shown in Fig. 12 is an example of one of the images that may be captured at time t from an environment of a host vehicle traveling in lane 1210 along a predicted trajectory 1212. The navigation system may include at least one processing device (e.g., including any of the EyeQ processors or other devices described above) that are specifically programmed to receive the plurality of images and analyze the images to determine an action in response to the scene. Specifically, the at least one processing device may implement sensing module 801, driving

policy module 803, and control module 805, as shown in Fig. 8. Sensing module 801 may be responsible for collecting and outputting the image information collected from the cameras and providing that information, in the form of an identified navigational state, to driving policy module 803, which may constitute a trained navigational system that has been trained through machine learning techniques, such as supervised learning, reinforcement learning, etc. Based on the navigational state information provided to driving policy module 803 by sensing module 801, driving policy module 803 (e.g., by implementing the options graph approach described above) may generate a desired navigational action for execution by the host vehicle in response to the identified navigational state.

[0295] In some embodiments, the at least one processing device may translate the desired navigation action directly into navigational commands using, for example, control module 805. In other embodiments, however, hard constraints may be applied such that the desired navigational action provided by the driving policy module 803 is tested against various predetermined navigational constraints that may be implicated by the scene and the desired navigational action. For example, where driving policy module 803 outputs a desired navigational action that would cause the host vehicle to follow trajectory 1212, this navigational action may be tested relative to one or more hard constraints associated with various aspects of the environment of the host vehicle. For example, a captured image 1201 may reveal a curb 1213, a pedestrian 1215, a target vehicle 1217, and a stationary object (e.g., an overturned box) present in the scene. Each of these may be associated with one or more hard constraints. For example, curb 1213 may be associated with a static constraint that prohibits the host vehicle from navigating into the curb or past the curb and onto a sidewalk 1214. Curb 1213 may also be associated with a road barrier envelope that defines a distance (e.g., a buffer zone) extending away from (e.g., by 0.1 meters, 0.25 meters, 0.5 meters, 1 meter, etc.) and along the curb, which defines a no-navigate zone for the host vehicle. Of course, static constraints may be associated with other types of roadside boundaries as well (e.g., guard rails, concrete pillars, traffic cones, pylons, or any other type of roadside barrier).

[0296] It should be noted that distances and ranging may be determined by any suitable method. For example, in some embodiments, distance information may be provided by onboard radar and/or lidar systems. Alternatively or additionally, distance information may be derived from analysis of one or more images captured from the environment of the host vehicle. For example, numbers of pixels of a recognized object represented in an image may be determined and compared to known field of view and focal length geometries of the image capture devices to determine scale and distances. Velocities and accelerations may be determined, for example, by observing changes in scale between objects from image to image over known time intervals. This analysis may indicate the direction of movement toward or away from the host vehicle along with how fast the object is pulling away from or coming toward the host vehicle. Crossing velocity may be determined through analysis of the change in an object's X coordinate position from one image to another over known time periods.

[0297] Pedestrian 1215 may be associated with a pedestrian envelope that defines a buffer zone 1216. In some cases, an imposed hard constraint may prohibit the host vehicle from navigating within a distance of 1 meter from pedestrian 1215 (in any direction relative to the pedestrian). Pedestrian

1215 may also define the location of a pedestrian influence zone 1220. Such an influence zone may be associated with a constraint that limits the speed of the host vehicle within the influence zone. The influence zone may extend 5 meters, 10 meters, 20 meters, etc., from pedestrian 1215. Each graduation of the influence zone may be associated with a different speed limit. For example, within a zone of 1
5 meter to five meters from pedestrian 1215, host vehicle may be limited to a first speed (e.g., 10 mph, 20 mph, etc.) that may be less than a speed limit in a pedestrian influence zone extending from 5 meters to 10 meters. Any graduation for the various stages of the influence zone may be used. In some
embodiments, the first stage may be narrower than from 1 meter to five meters and may extend only from one meter to two meters. In other embodiments, the first stage of the influence zone may extend from 1
10 meter (the boundary of the no-navigate zone around a pedestrian) to a distance of at least 10 meters. A second stage, in turn, may extend from 10 meters to at least about 20 meters. The second stage may be associated with a maximum rate of travel for the host vehicle that is greater than the maximum rate of travel associated with the first stage of the pedestrian influence zone.

[0298] One or more stationary object constraints may also be implicated by the detected scene
15 in the environment of the host vehicle. For example, in image 1201, the at least one processing device may detect a stationary object, such as box 1219 present in the roadway. Detected stationary objects may include various objects, such as at least one of a tree, a pole, a road sign, or an object in a roadway. One or more predefined navigational constraints may be associated with the detected stationary object. For example, such constraints may include a stationary object envelope, wherein the stationary object
20 envelope defines a buffer zone about the object within which navigation of the host vehicle may be prohibited. At least a portion of the buffer zone may extend a predetermined distance from an edge of the detected stationary object. For example, in the scene represented by image 1201, a buffer zone of at least 0.1 meters, 0.25 meters, 0.5 meters or more may be associated with box 1219 such that the host vehicle will pass to the right or to the left of the box by at least some distance (e.g., the buffer zone distance) in
25 order to avoid a collision with the detected stationary object.

[0299] The predefined hard constraints may also include one or more target vehicle constraints. For example, a target vehicle 1217 may be detected in image 1201. To ensure that the host vehicle does not collide with target vehicle 1217, one or more hard constraints may be employed. In some cases, a target vehicle envelope may be associated with a single buffer zone distance. For example,
30 the buffer zone may be defined by a 1 meter distance surrounding the target vehicle in all directions. The buffer zone may define a region extending from the target vehicle by at least one meter into which the host vehicle is prohibited from navigating.

[0300] The envelope surrounding target vehicle 1217 need not be defined by a fixed buffer distance, however. In some cases the predefined hard constraints associate with target vehicles (or any
35 other movable objects detected in the environment of the host vehicle) may depend on the orientation of the host vehicle relative to the detected target vehicle. For example, in some cases, a longitudinal buffer zone distance (e.g., one extending from the target vehicle toward the front or rear of the host vehicle—such as in the case that the host vehicle is driving toward the target vehicle) may be at least one meter. A

lateral buffer zone distance (e.g., one extending from the target vehicle toward either side of the host vehicle—such as when the host vehicle is traveling in a same or opposite direction as the target vehicle such that a side of the host vehicle will pass adjacent to a side of the target vehicle) may be at least 0.5 meters.

5 [0301] As described above, other constraints may also be implicated by detection of a target vehicle or a pedestrian in the environment of the host vehicle. For example, the predicted trajectories of the host vehicle and target vehicle 1217 may be considered and where the two trajectories intersect (e.g., at intersection point 1230), a hard constraint may require $t_1^l > t_2^a$ or $t_1^l > t_2^a + 0.5$ where the host vehicle is vehicle 1, and target vehicle 1217 is vehicle 2. Similarly, the trajectory of pedestrian 1215
10 (based on a heading direction and speed) may be monitored relative to the projected trajectory of the host vehicle. Given a particular pedestrian trajectory, with every point p on the trajectory, t(p) will represent the time required for the pedestrian to reach point p (i.e., point 1231 in Fig. 12). To maintain the required buffer distance of at least 1 meter from the pedestrian, either t(p) must be larger than the time the host vehicle will reach point p (with sufficient difference in time such that the host vehicle passes in front of
15 the pedestrian by a distance of at least one meter) or that t(p) must be less than the time the host vehicle will reach point p (e.g., if the host vehicle brakes to give way to the pedestrian). Still, in the latter example, the hard constraint will require that the host vehicle arrive at point p at a sufficient time later than the pedestrian such that the host vehicle can pass behind the pedestrian and maintain the required buffer distance of at least one meter.

20 [0302] Other hard constraints may also be employed. For example, a maximum deceleration rate of the host vehicle may be employed in at least some cases. Such a maximum deceleration rate may be determined based on a detected distance to a target vehicle following the host vehicle (e.g., using images collected from a rearward facing camera). The hard constraints may include a mandatory stop at a sensed crosswalk or a railroad crossing or other applicable constraints.

25 [0303] Where analysis of a scene in an environment of the host vehicle indicates that one or more predefined navigational constraints may be implicated, those constraints may be imposed relative to one or more planned navigational actions for the host vehicle. For example, where analysis of a scene results in driving policy module 803 returning a desired navigational action, that desired navigational action may be tested against one or more implicated constraints. If the desired navigational action is
30 determined to violate any aspect of the implicated constraints (e.g., if the desired navigational action would carry the host vehicle within a distance of 0.7 meters of pedestrian 1215 where a predefined hard constraint requires that the host vehicle remain at least 1.0 meters from pedestrian 1215), then at least one modification to the desired navigational action may be made based on the one or more predefined navigational constraints. Adjusting the desired navigational action in this way may provide an actual
35 navigational action for the host vehicle in compliance with the constraints implicated by a particular scene detected in the environment of the host vehicle.

[0304] After determination of the actual navigational action for the host vehicle, that navigational action may be implemented by causing at least one adjustment of a navigational actuator of

the host vehicle in response to the determined actual navigational action for the host vehicle. Such navigational actuator may include at least one of a steering mechanism, a brake, or an accelerator of the host vehicle.

[0305] Prioritized Constraints

5 [0306] As described above, various hard constraints may be employed with a navigational system to ensure safe operation of a host vehicle. The constraints may include a minimum safe driving distance with respect to a pedestrian, a target vehicle, a road barrier, or a detected object, a maximum speed of travel when passing within an influence zone of a detected pedestrian, or a maximum deceleration rate for the host vehicle, among others. These constraints may be imposed with a trained
10 system trained based on machine learning (supervised, reinforcement, or a combination), but they also may be useful with non-trained systems (e.g., those employing algorithms to directly address anticipated situations arising in scenes from a host vehicle environment).

[0307] In either case, there may be a hierarchy of constraints. In other words, some navigational constraints may have priority over other constraints. Thus, if a situation arose in which a
15 navigational action was not available that would result in all implicated constraints being satisfied, the navigation system may determine the available navigational action that achieves the highest priority constraints first. For example, the system may cause the vehicle to avoid a pedestrian first even if navigation to avoid the pedestrian would result in a collision with another vehicle or an object detected in a road. In another example, the system may cause the vehicle to ride up on a curb to avoid a pedestrian.

20 [0308] Fig. 13 provides a flowchart illustrating an algorithm for implementing a hierarchy of implicated constraints determined based on analysis of a scene in an environment of a host vehicle. For example, at step 1301, at least one processing device associated with the navigational system (e.g., an EyeQ processor, etc.) may receive, from a camera mounted on the host vehicle, a plurality of images representative of an environment of the host vehicle. Through analysis of an image or images
25 representative of the scene of the host vehicle environment at step 1303, a navigational state associated with the host vehicle may be identified. For example, a navigational state may indicate that the host vehicle is traveling along a two-lane road 1210, as in Fig. 12, that a target vehicle 1217 is moving through an intersection ahead of the host vehicle, that a pedestrian 1215 is waiting to cross the road on which the host vehicle travels, that an object 1219 is present ahead in the host vehicle lane, among various other
30 attributes of the scene.

[0309] At step 1305, one or more navigational constraints implicated by the navigational state of the host vehicle may be determined. For example, the at least one processing device, after analyzing a scene in the environment of the host vehicle represented by one or more captured images may determine one or more navigational constraints implicated by objects, vehicles, pedestrians, etc., recognized through
35 image analysis of the captured images. In some embodiments, the at least one processing device may determine at least a first predefined navigational constraint and a second predefined navigational constraint implicated by the navigational state, and the first predefined navigational constraint may differ from the second predefined navigational constraint. For example, the first navigational constraint may

relate to one or more target vehicles detected in the environment of the host vehicle, and the second navigational constraint may relate to a pedestrian detected in the environment of the host vehicle.

[0310] At step 1307, the at least one processing device may determine a priority associated with constraints identified in step 1305. In the example described, the second predefined navigational constraint, relating to pedestrians, may have a priority higher than the first predefined navigational constraint, which relates to target vehicles. While priorities associated with navigational constraints may be determined or assigned based on various factors, in some embodiments, the priority of a navigational constraint may be related to its relative importance from a safety perspective. For example, while it may be important that all implemented navigational constraints be followed or satisfied in as many situations as possible, some constraints may be associated with greater safety risks than others and, therefore, may be assigned higher priorities. For example, a navigational constraint requiring that the host vehicle maintain at least a 1 meter spacing from a pedestrian may have a higher priority than a constraint requiring that the host vehicle maintain at least a 1 meter spacing from a target vehicle. This may be because a collision with a pedestrian may have more severe consequences than a collision with another vehicle. Similarly, maintaining a space between the host vehicle and a target vehicle may have a higher priority than a constraint requiring the host vehicle to avoid a box in the road, to drive less than a certain speed over a speed bump, or to expose the host vehicle occupants to no more than a maximum acceleration level.

[0311] While driving policy module 803 is designed to maximize safety by satisfying navigational constraints implicated by a particular scene or navigational state, in some situations it may be physically impossible to satisfy every implicated constraint. In such situations, the priority of each implicated constraint may be used to determine which of the implicated constraints should be satisfied first, as shown at step 1309. Continuing with the example above, in a situation where it is not possible satisfy both the pedestrian gap constraint and the target vehicle gap constraint, but rather only one of the constraints can be satisfied, then the higher priority of the pedestrian gap constraint may result in that constraint being satisfied before attempting to maintain a gap to the target vehicle. Thus, in normal situations, the at least one processing device may determine, based on the identified navigational state of the host vehicle, a first navigational action for the host vehicle satisfying both the first predefined navigational constraint and the second predefined navigational constraint where both the first predefined navigational constraint and the second predefined navigational constraint can be satisfied, as shown at step 1311. In other situations, however, where not all the implicated constraints can be satisfied, the at least one processing device may determine, based on the identified navigational state, a second navigational action for the host vehicle satisfying the second predefined navigational constraint (i.e., the higher priority constraint), but not satisfying the first predefined navigational constraint (having a priority lower than the second navigational constraint), where the first predefined navigational constraint and the second predefined navigational constraint cannot both be satisfied, as shown at step 1313.

[0312] Next, at step 1315, to implement the determined navigational actions for the host vehicle the at least one processing device can cause at least one adjustment of a navigational actuator of

the host vehicle in response to the determined first navigational action or the determined second navigational action for the host vehicle. As in previous example, the navigational actuator may include at least one of a steering mechanism, a brake, or an accelerator.

[0313] Constraint Relaxation

5 [0314] As discussed above, navigational constraints may be imposed for safety purposes. The constraints may include a minimum safe driving distance with respect to a pedestrian, a target vehicle, a road barrier, or a detected object, a maximum speed of travel when passing within an influence zone of a detected pedestrian, or a maximum deceleration rate for the host vehicle, among others. These constraints may be imposed in a learning or non-learning navigational system. In certain situations, these constraints
10 may be relaxed. For example, where the host vehicle slows or stops near a pedestrian, then progresses slowly to convey an intention to pass by the pedestrian, a response of the pedestrian can be detected from acquired images. If the response of the pedestrian is to stay still or to stop moving (and/or if eye contact with the pedestrian is sensed), it may be understood that the pedestrian recognizes an intent of the navigational system to pass by the pedestrian. In such situations, the system may relax one or more
15 predefined constraints and implement a less stringent constraint (e.g., allow the vehicle to navigate within 0.5 meters of a pedestrian rather than within a more stringent 1 meter boundary).

[0315] Fig. 14 provides a flowchart for implementing control of the host vehicle based on relaxation of one or more navigational constraints. At step 1401, the at least one processing device may receive, from a camera associated with the host vehicle, a plurality of images representative of an
20 environment of the host vehicle. Analysis of the images at step 1403 may enable identification of a navigational state associated with the host vehicle. At step 1405, the at least one processor may determine navigational constraints associated with the navigational state of the host vehicle. The navigational constraints may include a first predefined navigational constraint implicated by at least one aspect of the navigational state. At step 1407, analysis of the plurality of images may reveal the presence of at least
25 one navigational constraint relaxation factor.

[0316] A navigational constraint relaxation factor may include any suitable indicator that one or more navigational constraints may be suspended, altered, or otherwise relaxed in at least one aspect. In some embodiments, the at least one navigational constraint relaxation factor may include a determination (based on image analysis) that the eyes of a pedestrian are looking in a direction of the host vehicle. In
30 such cases, it may more safely be assumed that the pedestrian is aware of the host vehicle. As a result, a confidence level may be higher that the pedestrian will not engage in unexpected actions that cause the pedestrian to move into a path of the host vehicle. Other constraint relaxation factors may also be used. For example, the at least one navigational constraint relaxation factor may include: a pedestrian determined to be not moving (e.g., one presumed to be less likely of entering a path of the host vehicle);
35 or a pedestrian whose motion is determined to be slowing. The navigational constraint relaxation factor may also include more complicated actions, such as a pedestrian determined to be not moving after the host vehicle has come to a stop and then resumed movement. In such a situation, the pedestrian may be assumed to understand that the host vehicle has a right of way, and the pedestrian coming to a stop may

suggest an intent of the pedestrian to give way to the host vehicle. Other situations that may cause one or more constraints to be relaxed include the type of curb stone (e.g., a low curb stone or one with a gradual slope might allow a relaxed distance constraint), lack of pedestrians or other objects on sidewalk, a vehicle with its engine not running may have a relaxed distance, or a situation in which a pedestrian is facing away and/or is moving away from the area towards which the host vehicle is heading.

[0317] Where the presence of a navigational constraint relaxation factor is identified (e.g., at step 1407), a second navigational constraint may be determined or developed in response to detection of the constraint relaxation factor. This second navigational constraint may be different from the first navigational constraint and may include at least one characteristic relaxed with respect to the first navigational constraint. The second navigational constraint may include a newly generated constraint based on the first constraint, where the newly generated constraint includes at least one modification that relaxes the first constraint in at least one respect. Alternatively, the second constraint may constitute a predetermined constraint that is less stringent than the first navigational constraint in at least one respect. In some embodiments, such second constraints may be reserved for usage only for situations where a constraint relaxation factor is identified in an environment of the host vehicle. Whether the second constraint is newly generated or selected from a set of fully or partially available predetermined constraints, application of a second navigational constraint in place of a more stringent first navigational constraint (that may be applied in the absence of detection of relevant navigational constraint relaxation factors) may be referred to as constraint relaxation and may be accomplished in step 1409.

[0318] Where at least one constraint relaxation factor is detected at step 1407, and at least one constraint has been relaxed in step 1409, a navigational action for the host vehicle may be determined at step 1411. The navigational action for the host vehicle may be based on the identified navigational state and may satisfy the second navigational constraint. The navigational action may be implemented at step 1413 by causing at least one adjustment of a navigational actuator of the host vehicle in response to the determined navigational action.

[0319] As discussed above, the usage of navigational constraints and relaxed navigational constraints may be employed with navigational systems that are trained (e.g., through machine learning) or untrained (e.g., systems programmed to respond with predetermined actions in response to specific navigational states). Where trained navigational systems are used, the availability of relaxed navigational constraints for certain navigational situations may represent a mode switching from a trained system response to an untrained system response. For example, a trained navigational network may determine an original navigational action for the host vehicle, based on the first navigational constraint. The action taken by the vehicle, however, may be one that is different from the navigational action satisfying the first navigational constraint. Rather, the action taken may satisfy the more relaxed second navigational constraint and may be an action developed by a non-trained system (e.g., as a response to detection of a particular condition in the environment of the host vehicle, such as the presence of a navigational constraint relaxation factor).

[0320] There are many examples of navigational constraints that may be relaxed in response to detection in the environment of the host vehicle of a constraint relaxation factor. For example, where a predefined navigational constraint includes a buffer zone associated with a detected pedestrian, and at least a portion of the buffer zone extends a distance from the detected pedestrian, a relaxed navigational constraint (either newly developed, called up from memory from a predetermined set, or generated as a relaxed version of a preexisting constraint) may include a different or modified buffer zone. For example, the different or modified buffer zone may have a distance relative to the pedestrian that is less than the original or unmodified buffer zone relative to the detected pedestrian. As a result, in view of the relaxed constraint, the host vehicle may be permitted to navigate closer to a detected pedestrian, where an appropriate constraint relaxation factor is detected in the environment of the host vehicle.

[0321] A relaxed characteristic of a navigational constraint may include a reduced width in a buffer zone associated with at least one pedestrian, as noted above. The relaxed characteristic, however, may also include a reduced width in a buffer zone associated with a target vehicle, a detected object, a roadside barrier, or any other object detected in the environment of the host vehicle.

[0322] The at least one relaxed characteristic may also include other types of modifications in navigational constraint characteristics. For example, the relaxed characteristic may include an increase in speed associated with at least one predefined navigational constraint. The relaxed characteristic may also include an increase in a maximum allowable deceleration/acceleration associated with at least one predefined navigational constraint.

[0323] While constraints may be relaxed in certain situations, as described above, in other situations, navigational constraints may be augmented. For example, in some situations, a navigational system may determine that conditions warrant augmentation of a normal set of navigational constraints. Such augmentation may include adding new constraints to a predefined set of constraints or adjusting one or more aspects of a predefined constraint. The addition or adjustment may result in more conservative navigation relative the predefined set of constraints applicable under normal driving conditions. Conditions that may warrant constraint augmentation may include sensor failure, adverse environmental conditions (rain, snow, fog, or other conditions associated with reduced visibility or reduced vehicle traction), etc.

[0324] Fig. 15 provides a flowchart for implementing control of the host vehicle based on augmentation of one or more navigational constraints. At step 1501, the at least one processing device may receive, from a camera associated with the host vehicle, a plurality of images representative of an environment of the host vehicle. Analysis of the images at step 1503 may enable identification of a navigational state associated with the host vehicle. At step 1505, the at least one processor may determine navigational constraints associated with the navigational state of the host vehicle. The navigational constraints may include a first predefined navigational constraint implicated by at least one aspect of the navigational state. At step 1507, analysis of the plurality of images may reveal the presence of at least one navigational constraint augmentation factor.

[0325] An implicated navigational constraint may include any of the navigational constraints discussed above (e.g., with respect to Fig. 12) or any other suitable navigational constraints. A navigational constraint augmentation factor may include any indicator that one or more navigational constraints may be supplemented/augmented in at least one aspect. Supplementation or augmentation of navigational constraints may be performed on a per set basis (e.g., by adding new navigational constraints to a predetermined set of constraints) or may be performed on a per constraint basis (e.g., modifying a particular constraint such that the modified constraint is more restrictive than the original, or adding a new constraint that corresponds to a predetermined constraint, wherein the new constraint is more restrictive than the corresponding constraint in at least one aspect). Additionally, or alternatively, supplementation or augmentation of navigational constraints may refer selection from among a set of predetermined constraints based on a hierarchy. For example, a set of augmented constraints may be available for selection based on whether a navigational augmentation factor is detected in the environment of or relative to the host vehicle. Under normal conditions where no augmentation factor is detected, then the implicated navigational constraints may be drawn from constraints applicable to normal conditions. On the other hand, where one or more constraint augmentation factors are detected, the implicated constraints may be drawn from augmented constraints either generated or predefined relative to the one or more augmentation factors. The augmented constraints may be more restrictive in at least one aspect than corresponding constraints applicable under normal conditions.

[0326] In some embodiments, the at least one navigational constraint augmentation factor may include a detection (e.g., based on image analysis) of the presence of ice, snow, or water on a surface of a road in the environment of the host vehicle. Such a determination may be based, for example, upon detection of: areas of reflectance higher than expected for dry roadways (e.g., indicative of ice or water on the roadway); white regions on the road indicating the presence of snow; shadows on the roadway consistent with the presence of longitudinal trenches (e.g., tire tracks in snow) on the roadway; water droplets or ice/snow particles on a windshield of the host vehicle; or any other suitable indicator of the presence of water or ice/snow on a surface of a road.

[0327] The at least one navigational constraint augmentation factor may also include detection of particulates on an outer surface of a windshield of the host vehicle. Such particulates may impair image quality of one or more image capture devices associated with the host vehicle. While described with respect to a windshield of the host vehicle, which is relevant for cameras mounted behind the windshield of the host vehicle, detection of particulates on other surfaces (e.g., a lens or lens cover of a camera, headlight lens, rear windshield, a tail light lens, or any other surface of the host vehicle visible to an image capture device (or detected by a sensor) associated with the host vehicle may also indicate the presence of a navigational constraint augmentation factor.

[0328] The navigational constraint augmentation factor may also be detected as an attribute of one or more image acquisition devices. For example, a detected decrease in image quality of one or more images captured by an image capture device (e.g., a camera) associated with the host vehicle may also constitute a navigational constraint augmentation factor. A decline in image quality may be associated

with a hardware failure or partial hardware failure associated with the image capture device or an assembly associated with the image capture device. Such a decline in image quality may also be caused by environmental conditions. For example, the presence of smoke, fog, rain, snow, etc., in the air surrounding the host vehicle may also contribute to reduced image quality relative to the road, pedestrians, target vehicles, etc., that may be present in an environment of the host vehicle.

[0329] The navigational constraint augmentation factor may also relate to other aspects of the host vehicle. For example, in some situations, the navigational constraint augmentation factor may include a detected failure or partial failure of a system or sensor associated with the host vehicle. Such an augmentation factor may include, for example, detection of failure or partial failure of a speed sensor, GPS receiver, accelerometer, camera, radar, lidar, brakes, tires, or any other system associated with the host vehicle that may impact the ability of the host vehicle to navigate relative to navigational constraints associated with a navigational state of the host vehicle.

[0330] Where the presence of a navigational constraint augmentation factor is identified (e.g., at step 1507), a second navigational constraint may be determined or developed in response to detection of the constraint augmentation factor. This second navigational constraint may be different from the first navigational constraint and may include at least one characteristic augmented with respect to the first navigational constraint. The second navigational constraint may be more restrictive than the first navigational constraint, because detection of a constraint augmentation factor in the environment of the host vehicle or associated with the host vehicle may suggest that the host vehicle may have at least one navigational capability reduced with respect to normal operating conditions. Such reduced capabilities may include lowered road traction (e.g., ice, snow, or water on a roadway; reduced tire pressure; etc.); impaired vision (e.g., rain, snow, dust, smoke, fog etc. that reduces captured image quality); impaired detection capability (e.g., sensor failure or partial failure, reduced sensor performance, etc.), or any other reduction in capability of the host vehicle to navigate in response to a detected navigational state.

[0331] Where at least one constraint augmentation factor is detected at step 1507, and at least one constraint has been augmented in step 1509, a navigational action for the host vehicle may be determined at step 1511. The navigational action for the host vehicle may be based on the identified navigational state and may satisfy the second navigational (i.e., augmented) constraint. The navigational action may be implemented at step 1513 by causing at least one adjustment of a navigational actuator of the host vehicle in response to the determined navigational action.

[0332] As discussed, the usage of navigational constraints and augmented navigational constraints may be employed with navigational systems that are trained (e.g., through machine learning) or untrained (e.g., systems programmed to respond with predetermined actions in response to specific navigational states). Where trained navigational systems are used, the availability of augmented navigational constraints for certain navigational situations may represent a mode switching from a trained system response to an untrained system response. For example, a trained navigational network may determine an original navigational action for the host vehicle, based on the first navigational constraint. The action taken by the vehicle, however, may be one that is different from the navigational action

satisfying the first navigational constraint. Rather, the action taken may satisfy the augmented second navigational constraint and may be an action developed by a non-trained system (e.g., as a response to detection of a particular condition in the environment of the host vehicle, such as the presence of a navigational constraint augmented factor).

5 [0333] There are many examples of navigational constraints that may be generated, supplemented, or augmented in response to detection in the environment of the host vehicle of a constraint augmentation factor. For example, where a predefined navigational constraint includes a buffer zone associated with a detected pedestrian, object, vehicle, etc., and at least a portion of the buffer zone extends a distance from the detected pedestrian/object/vehicle, an augmented navigational constraint
10 (either newly developed, called up from memory from a predetermined set, or generated as an augmented version of a preexisting constraint) may include a different or modified buffer zone. For example, the different or modified buffer zone may have a distance relative to the pedestrian/object/vehicle that is greater than the original or unmodified buffer zone relative to the detected pedestrian/object/vehicle. As a result, in view of the augmented constraint, the host vehicle may be forced to navigate further from the
15 detected pedestrian/object/vehicle, where an appropriate constraint augmentation factor is detected in the environment of the host vehicle or relative to the host vehicle.

[0334] The at least one augmented characteristic may also include other types of modifications in navigational constraint characteristics. For example, the augmented characteristic may include a decrease in speed associated with at least one predefined navigational constraint. The augmented
20 characteristic may also include a decrease in a maximum allowable deceleration/acceleration associated with at least one predefined navigational constraint.

[0335] Navigation Based on Long Range Planning

[0336] In some embodiments, the disclosed navigational system can respond not only to a detected navigational state in an environment of the host vehicle, but may also determine one or more
25 navigational actions based on long range planning. For example, the system may consider the potential impact on future navigational states of one or more navigational actions available as options for navigating with respect to a detected navigational state. Considering the effects of available actions on future states may enable the navigational system to determine navigational actions based not just upon a currently detected navigational state, but also based upon long range planning. Navigation using long
30 range planning techniques may be especially applicable where one or more reward functions are employed by the navigation system as a technique for selecting navigational actions from among available options. Potential rewards may be analyzed with respect to the available navigational actions that may be taken in response to a detected, current navigational state of the host vehicle. Further, however, the potential rewards may also be analyzed relative to actions that may be taken in response to
35 future navigational states projected to result from the available actions to a current navigational state. As a result, the disclosed navigational system may, in some cases, select a navigational action in response to a detected navigational state even where the selected navigational action may not yield the highest reward from among the available actions that may be taken in response to the current navigational state. This

may be especially true where the system determines that the selected action may result in a future navigational state giving rise to one or more potential navigational actions offering higher rewards than the selected action or, in some cases, any of the actions available relative to a current navigational state. The principle may be expressed more simply as taking a less favorable action now in order to produce higher reward options in the future. Thus, the disclosed navigational system capable of long range planning may choose a suboptimal short term action where long term prediction indicates that a short term loss in reward may result in long term reward gains.

[0337] In general, autonomous driving applications may involve a series of planning problems, where the navigational system may decide on immediate actions in order to optimize a longer term objective. For example, when a vehicle is confronted with a merge situation at a roundabout, the navigational system may decide on an immediate acceleration or braking command in order to initiate navigation into the roundabout. While the immediate action to the detected navigational state at the roundabout may involve an acceleration or braking command responsive to the detected state, the long term objective is a successful merge, and the long term effect of the selected command is the success/failure of the merge. The planning problem may be addressed by decomposing the problem into two phases. First, supervised learning may be applied for predicting the near future based on the present (assuming the predictor will be differentiable with respect to the representation of the present). Second, a full trajectory of the agent may be modeled using a recurrent neural network, where unexplained factors are modeled as (additive) input nodes. This may allow solutions to the long-term planning problem to be determined using supervised learning techniques and direct optimization over the recurrent neural network. Such an approach may also enable the learning of robust policies by incorporating adversarial elements to the environment.

[0338] Two of the most fundamental elements of autonomous driving systems are sensing and planning. Sensing deals with finding a compact representation of the present state of the environment, while planning deals with deciding on what actions to take so as to optimize future objectives. Supervised machine learning techniques are useful for solving sensing problems. Machine learning algorithmic frameworks may also be used for the planning part, especially reinforcement learning (RL) frameworks, such as those described above.

[0339] RL may be performed in a sequence of consecutive rounds. At round t , the planner (a.k.a. the agent or driving policy module 803) may observe a state, $s_t \in \mathcal{S}$, which represents the agent as well as the environment. It then should decide on an action $a_t \in \mathcal{A}$. After performing the action, the agent receives an immediate reward, $r_t \in \mathbb{R}$, and is moved to a new state, s_{t+1} . As an example, the host vehicle may include an adaptive cruise control (ACC) system, in which the vehicle should autonomously implement acceleration/braking so as to keep an adequate distance to a preceding vehicle while maintaining smooth driving. The state can be modeled as a pair, $s_t = (x_t, v_t) \in \mathbb{R}^2$, where x_t is the distance to the preceding vehicle and v_t is the velocity of the host vehicle relative to the velocity of the preceding vehicle. The action $a_t \in \mathbb{R}$ will be the acceleration command (where the host vehicle slows

down if $a_t < 0$). The reward can be a function that depends on $|a_t|$ (reflecting the smoothness of driving) and on s_t (reflecting that the host vehicle maintains a safe distance from the preceding vehicle). The goal of the planner is to maximize the cumulative reward (maybe up to a time horizon or a discounted sum of future rewards). To do so, the planner may rely on a policy, $\pi : S \rightarrow A$, which maps a state into an

5 action.

[0340] Supervised Learning (SL) can be viewed as a special case of RL, in which s_t is sampled from some distribution over S , and the reward function may have the form $r_t = -\ell(a_t, y_t)$, where ℓ is a loss function, and the learner observes the value of y_t which is the (possibly noisy) value of the optimal action to take when viewing the state s_t . There may be several differences between a general RL

10 model and a specific case of SL, and these differences can make the general RL problem more challenging.

[0341] In some SL situations, the actions (or predictions) taken by the learner may have no effect on the environment. In other words, s_{t+1} and a_t are independent. This can have two important implications. First, in SL, a sample $(s_1, y_1), \dots, (s_m, y_m)$ can be collected in advance, and only then

15 can the search begin for a policy (or predictor) that will have good accuracy relative to the sample. In contrast, in RL, the state s_{t+1} usually depends on the action taken (and also on the previous state), which in turn depends on the policy used to generate the action. This ties the data generation process to the policy learning process. Second, because actions do not affect the environment in SL, the contribution of the choice of a_t to the performance of π is local. Specifically, a_t only affects the value of the immediate

20 reward. In contrast, in RL, actions that are taken at round t might have a long-term effect on the reward values in future rounds.

[0342] In SL, the knowledge of the “correct” answer, y_t , together with the shape of the reward, $r_t = -\ell(a_t, y_t)$ may provide full knowledge of the reward for all possible choices of a_t , which may enable calculation of the derivative of the reward with respect to a_t . In contrast, in RL, a “one-shot” value

25 of the reward may be all that can be observed for a specific choice of action taken. This may be referred to as a “bandit” feedback. This is one of the most significant reasons for the need of “exploration” as a part of long term navigational planning, because in RL-based systems, if only “bandit” feedback is available, the system may not always know if the action taken was the best action to take.

[0343] Many RL algorithms rely, at least in part, on the mathematically elegant model of a

30 Markov Decision Process (MDP). The Markovian assumption is that the distribution of s_{t+1} is fully determined given s_t and a_t . This yields a closed form expression for the cumulative reward of a given policy in terms of the stationary distribution over states of the MDP. The stationary distribution of a policy can be expressed as a solution to a linear programming problem. This yields two families of algorithms: 1) optimization with respect to the primal problem, which may be referred to as policy search,

35 and 2) optimization with respect to a dual problem, whose variables are called the *value function*, V^π . The value function determines the expected cumulative reward if the MDP begins from the initial state, s , and from there actions are chosen according to π . A related quantity is the state-action value function,

$Q^\pi(s, a)$, which determines the cumulative reward assuming a start from state, s , an immediately chosen action a , and from there on actions chosen according to π . The Q function may give rise to a characterization of the optimal policy (using the Bellman's equation). In particular, the Q function may show that the optimal policy is a deterministic function from S to A (in fact, it may be characterized as a
5 "greedy" policy with respect to the optimal Q function).

[0344] One potential advantage of the MDP model is that it allows coupling of the future into the present using the Q function. For example, given that a host vehicle is now in state, s , the value of $Q^\pi(s, a)$ may indicate the effect of performing action a on the future. Therefore, the Q function may provide a local measure of the quality of an action a , thus making the RL problem more similar to a SL
10 scenario.

[0345] Many RL algorithms approximate the V function or the Q function in one way or another. Value iteration algorithms, e.g., the Q learning algorithm, may rely on the fact that the V and Q functions of the optimal policy may be fixed points of some operators derived from Bellman's equation. Actor-critic policy iteration algorithms aim to learn a policy in an iterative way, where at iteration t , the
15 "critic" estimates Q^{π_t} and based on this estimate, the "actor" improves the policy.

[0346] Despite the mathematical elegance of MDPs and the convenience of switching to the Q function representation, this approach may have several limitations. For example, an approximate notion of a Markovian behaving state may be all that can be found in some cases. Furthermore, the transition of states may depend not only on the agent's action, but also on actions of other players in the environment.
20 For example, in the ACC example mentioned above, while the dynamic of the autonomous vehicle may be Markovian, the next state may depend on the behavior of the driver of the other car, which is not necessarily Markovian. One possible solution to this problem is to use partially observed MDPs, in which it is assumed that there is a Markovian state, but an observation that is distributed according to the hidden state is what can be seen.

[0347] A more direct approach may consider game theoretical generalizations of MDPs (e.g., the Stochastic Games framework). Indeed, algorithms for MDPs may be generalized to multi-agents games (e.g., minimax- Q learning or Nash- Q learning). Other approaches may include explicit modeling of the other players and vanishing regret learning algorithms. Learning in a multi-agent setting may be
25 more complex than in a single agent setting.

[0348] A second limitation of the Q function representation may arise by departing from a tabular setting. The tabular setting is when the number of states and actions is small, and therefore, Q can be expressed as a table with $|S|$ rows and $|A|$ columns. However, if the natural representation of S and A includes Euclidean spaces, and the state and action spaces are discretized, the number of states/actions may be exponential in the dimension. In such cases, it may not be practical to employ a tabular setting.
30 Instead, the Q function may be approximated by some function from a parametric hypothesis class (e.g., neural networks of a certain architecture). For example, a deep- Q -network (DQN) learning algorithm may be used. In DQN, the state space can be continuous, but the action space may remain a small discrete set.
35

There may be approaches for dealing with continuous action spaces, but they may rely on approximating the Q function. In any case, the Q function may be complicated and sensitive to noise, and, therefore, may be challenging to learn.

[0349] A different approach may be to address the RL problem using a recurrent neural network (RNN). In some cases, RNN may be combined with the notions of multi-agents games and robustness to adversarial environments from game theory. Further, this approach may be one that does not explicitly rely on any Markovian assumption.

[0350] The following describes in more detail an approach for navigation by planning based on prediction. In this approach, it may be assumed that the state space, S , is a subset of \mathbb{R}^d , and the action space, A , is a subset of \mathbb{R}^k . This may be a natural representation in many applications. As noted above, there may be two key differences between RL and SL: (1) because past actions affect future rewards, information from the future may need to be propagated back to the past; and (2) the “bandit” nature of rewards can blur the dependence between (state, action) and reward, which can complicate the learning process.

[0351] As a first step in the approach, an observation may be made that there are interesting problems in which the bandit nature of rewards is not an issue. For example, reward value (as will be discussed in more detail below) for the ACC application may be differentiable with respect to the current state and action. In fact, even if the reward is given in a “bandit” manner, the problem of learning a differentiable function, $\hat{r}(s, a)$, such that $\hat{r}(s_t, a_t) \approx r_t$, may be a relatively straightforward SL problem (e.g., a one dimensional regression problem). Therefore, the first step of the approach may be to either define the reward as a function, $\hat{r}(s, a)$, which is differentiable with respect to s and a , or to use a regression learning algorithm in order to learn a differentiable function, \hat{r} , that minimizes at least some regression loss over a sample with instance vector being $(s_t, a_t) \in \mathbb{R}^d \times \mathbb{R}^k$ and target scalar being r_t . In some situations, in order to create a training set, elements of exploration may be used.

[0352] To address the connection between past and future, a similar idea may be used. For example, suppose a differentiable function $\hat{N}(s, a)$ can be learned such that $\hat{N}(s_t, a_t) \approx s_{t+1}$. Learning such a function may be characterized as an SL problem. \hat{N} may be viewed as a predictor for the near future. Next, a policy that maps from S to A may be described using a parametric function $\pi_\theta : S \rightarrow A$. Expressing π_θ as a neural network, may enable expression of an episode of running the agent for T rounds using a recurrent neural network (RNN), where the next state is defined as $s_{t+1} = \hat{N}(s_t, a_t) + \nu_t$. Here, $\nu_t \in \mathbb{R}^d$ may be defined by the environment and may express unpredictable aspects of the near future. The fact that s_{t+1} depends on s_t and a_t in a differentiable manner may enable a connection between future reward values and past actions. A parameter vector of the policy function, π_θ , may be learned by back-propagation over the resulting RNN. Note that explicit probabilistic assumptions need not be imposed on ν_t . In particular, there need not be a requirement for a

Markovian relation. Instead, the recurrent network may be relied upon to propagate “enough” information between past and future. Intuitively, $\hat{N}(s_t, a_t)$ may describe the predictable part of the near future, while ν_t may express the unpredictable aspects, which may arise due to the behavior of other players in the environment. The learning system should learn a policy that will be robust to the behavior of other

5 players. If $\|\nu_t\|$ is large, the connection between past actions and future reward values may be too noisy for learning a meaningful policy. Explicitly expressing the dynamic of the system in a transparent way may enable incorporation of prior knowledge more easily. For example, prior knowledge may simplify the problem of defining \hat{N} .

[0353] As discussed above, the learning system may benefit from robustness relative to an

10 adversarial environment, such as the environment of a host vehicle, which may include multiple other drivers that may act in unexpected way. In a model that does not impose probabilistic assumptions on ν_t , environments may be considered in which ν_t is chosen in an adversarial manner. In some cases, restrictions may be placed on μ_t , otherwise the adversary can make the planning problem difficult or even impossible. One natural restriction may be to require that $\|\mu_t\|$ is bounded by a constant.

15 [0354] Robustness against adversarial environments may be useful in autonomous driving applications. Choosing μ_t in an adversarial way may even speed up the learning process, as it can focus the learning system toward a robust optimal policy. A simple game may be used to illustrate this concept. The state is $s_t \in \mathbb{R}$, the action is $a_t \in \mathbb{R}$, and the immediate loss function is $0.1|a_t| + [|s_t| - 2]_+$, where $[x]_+ = \max\{x, 0\}$ is the ReLU (rectified linear unit) function. The next state is

20 $s_{t+1} = s_t + a_t + \nu_t$, where $\nu_t \in [-0.5, 0.5]$ is chosen for the environment in an adversarial manner. Here, the optimal policy may be written as a two layer network with ReLU: $a_t = -[s_t - 1.5]_+ + [-s_t - 1.5]_+$. Observe that when $|s_t| \in (1.5, 2]$, the optimal action may have a larger immediate loss than the action $a = 0$. Therefore, the system may plan for the future and may not rely solely on the immediate loss. Observe that the derivative of the loss with respect to a_t is $0.1 \text{sign}(a_t)$, and the

25 derivative with respect to s_t is $\mathbb{1}[|s_t| > 2] \text{sign}(s_t)$. In a situation in which $s_t \in (1.5, 2]$, the adversarial choice of ν_t would be to set $\nu_t = 0.5$ and, therefore, there may be a non-zero loss on round $t + 1$, whenever $a_t > 1.5 - s_t$. In such cases, the derivative of the loss may back-propagate directly to a_t . Thus, the adversarial choice of ν_t may help the navigational system obtain a non-zero back-propagation message in cases for which the choice of a_t is sub-optimal. Such a relationship may aid the

30 navigational system in selecting present actions based on an expectation that such a present action (even if that action would result in a suboptimal reward or even a loss) will provide opportunities in the future for more optimal actions that result in higher rewards.

[0355] Such an approach may be applied to virtually any navigational situation that may arise. The following describes the approach applied to one example: adaptive cruise control (ACC). In the

35 ACC problem, the host vehicle may attempt to maintain an adequate distance to a target vehicle ahead

(e.g., 1.5 seconds to the target car). Another goal may be to drive as smooth as possible while maintaining the desired gap. A model representing this situation may be defined as follows. The state space is \mathbb{R}^3 , and the action space is \mathbb{R} . The first coordinate of the state is the speed of the target car, the second coordinate is the speed of the host vehicle, and the last coordinate is the distance between the host vehicle and target vehicle (e.g., location of the host vehicle minus the location of the target along the road curve). The action to be taken by the host vehicle is the acceleration, and may be denoted by a_t . The quantity τ may denote the difference in time between consecutive rounds. While τ may be set to any suitable quantity, in one example, τ may be 0.1 seconds. Position, s_t , may be denoted as

10 $s_t = (v_t^{\text{target}}, v_t^{\text{host}}, x_t)$, and the (unknown) acceleration of the target vehicle may be denoted as a_t^{target} .

[0356] The full dynamics of the system can be described by:

$$\begin{aligned} v_t^{\text{target}} &= [v_{t-1}^{\text{target}} + \tau a_{t-1}^{\text{target}}]_+ \\ v_t^{\text{host}} &= [v_{t-1}^{\text{host}} + \tau a_{t-1}]_+ \\ x_t &= [x_{t-1} + \tau (v_{t-1}^{\text{target}} - v_{t-1}^{\text{host}})]_+ \end{aligned}$$

[0357] This can be described as a sum of two vectors:

$$\begin{aligned} s_t &= ([s_{t-1}[0] + \tau a_{t-1}^{\text{target}}]_+, [s_{t-1}[1] + \tau a_{t-1}]_+, [s_{t-1}[2] + \tau (s_{t-1}[0] - s_{t-1}[1])]_+) \\ &= \underbrace{([s_{t-1}[0], [s_{t-1}[1] + \tau a_{t-1}]_+, [s_{t-1}[2] + \tau (s_{t-1}[0] - s_{t-1}[1])]_+])}_{\tilde{N}(s_{t-1}, a_t)} + \underbrace{([s_{t-1}[0] + \tau a_{t-1}^{\text{target}}]_+ - s_{t-1}[0], 0, 0)}_{v_t} \end{aligned}$$

15 [0358] The first vector is the predictable part, and the second vector is the unpredictable part. The reward on round t is defined as follows:

$$-r_t = 0.1 |a_t| + [|x_t/x_t^* - 1| - 0.3]_+ \quad \text{where} \quad x_t^* = \max\{1, 1.5 v_t^{\text{host}}\}$$

[0359] The first term may result in a penalty for non-zero accelerations, thus encouraging smooth driving. The second term depends on the ratio between the distance to the target car, x_t , and the desired distance, x_t^* , which is defined as the maximum between a distance of 1 meter and break distance of 1.5 seconds. In some cases, this ratio may be exactly 1, but as long as this ratio is within [0.7, 1.3], the policy may forego any penalties, which may allow the host vehicle some slack in navigation—a characteristic that may be important in achieving a smooth drive.

25 [0360] Implementing the approach outlined above, the navigation system of the host vehicle (e.g., through operation of driving policy module 803 within processing unit 110 of the navigation system) may select an action in response to an observed state. The selected action may be based on analysis not only of rewards associated with the responsive actions available relative to a sensed navigational state, but may also be based on consideration and analysis of future states, potential actions in response to the futures states, and rewards associated with the potential actions.

30 [0361] Fig. 16 illustrates an algorithmic approach to navigation based on detection and long range planning. For example, at step 1601, the at least one processing device 110 of the navigation system for the host vehicle may receive a plurality of images. These images may capture scenes

representative of an environment of the host vehicle and may be supplied by any of the image capture devices (e.g., cameras, sensors, etc.) described above. Analysis of one or more of these images at step 1603 may enable the at least one processing device 110 to identify a present navigational state associated with the host vehicle (as described above).

5 [0362] At steps 1605, 1607, and 1609, various potential navigational actions responsive to the sensed navigational state may be determined. These potential navigational actions (e.g., a first navigational action through an Nth available navigational action) may be determined based on the sensed state and the long range goals of the navigational system (e.g., to complete a merge, follow a lead vehicle smoothly, pass a target vehicle, avoid an object in the roadway, slow for a detected stop sign, avoid a
10 target vehicle cutting in, or any other navigational action that may advance the navigational goals of the system).

[0363] For each of the determined potential navigational actions, the system may determine an expected reward. The expected reward may be determined according to any of the techniques described above and may include analysis of a particular potential action relative to one or more reward functions.
15 Expected rewards 1606, 1608, and 1610 may be determined for each of the potential navigational actions (e.g., the first, second, and Nth) determined in steps 1605, 1607, and 1609, respectively.

[0364] In some cases, the navigational system of the host vehicle may select from among the available potential actions based on values associated with expected rewards 1606, 1608, and 1610 (or any other type of indicator of an expected reward). For example, in some situations, the action that yields
20 the highest expected reward may be selected.

[0365] In other cases, especially where the navigation system engages in long range planning to determine navigational actions for the host vehicle, the system may not choose the potential action that yields the highest expected reward. Rather, the system may look to the future to analyze whether there may be opportunities for realizing higher rewards later if lower reward actions are selected in response to
25 a current navigational state. For example, for any or all of the potential actions determined at steps 1605, 1607, and 1609, a future state may be determined. Each future state, determined at steps 1613, 1615, and 1617, may represent a future navigational state expected to result based on the current navigational state as modified by a respective potential action (e.g., the potential actions determined at steps 1605, 1607, and 1609).

30 [0366] For each of the future states predicted at steps 1613, 1615, and 1617, one or more future actions (as navigational options available in response to determined future state) may be determined and evaluated. At steps 1619, 1621, and 1623, for example, values or any other type of indicator of expected rewards associated with one or more of the future actions may be developed (e.g., based on one or more reward functions). The expected rewards associated with the one or more future
35 actions may be evaluated by comparing values of reward functions associated with each future action or by comparing any other indicators associated with the expected rewards.

[0367] At step 1625, the navigational system for the host vehicle may select a navigational action for the host vehicle based on a comparison of expected rewards, not just based on the potential

actions identified relative to a current navigational state (e.g., at steps 1605, 1607, and 1609), but also based on expected rewards determined as a result of potential future actions available in response to predicted future states (e.g., determined at steps 1613, 1615, and 1617). The selection at step 1625 may be based on the options and rewards analysis performed at steps 1619, 1621, and 1623.

5 [0368] The selection of a navigational action at step 1625 may be based on a comparison of expected rewards associated with future action options only. In such a case, the navigational system may select an action to the current state based solely on a comparison of expected rewards resulting from actions to potential future navigational states. For example, the system may select the potential action identified at step 1605, 1607, or 1609 that is associated with a highest future reward value as determined
10 through analysis at steps 1619, 1621, and 1623.

[0369] The selection of a navigational action at step 1625 may also be based on comparison of current action options only (as noted above). In this situation, the navigational system may select the potential action identified at step 1605, 1607, or 1609 that is associated with a highest expected reward, 1606, 1608, or 1610. Such a selection may be performed with little or no consideration of future
15 navigational states or future expected rewards to navigational actions available in response to expected future navigational states.

[0370] On the other hand, in some cases, the selection of a navigational action at step 1625 may be based on a comparison of expected rewards associated with both future action options and with current action options. This, in fact, may be one of the principles of navigation based on long range
20 planning. For example, expected rewards to future actions may be analyzed to determine if any may warrant a selection of a lower reward action in response to the current navigational state in order to achieve a potential higher reward in response to a subsequent navigational action expected to be available in response to future navigational states. As an example, a value or other indicator of an expected reward 1606 may indicate a highest expected reward from among rewards 1606, 1608, and 1610. On the other
25 hand, expected reward 1608 may indicate a lowest expected reward from among rewards 1606, 1608, and 1610. Rather than simply selecting the potential action determined at step 1605 (i.e., the action giving rise to the highest expected reward 1606), analysis of future states, potential future actions, and future rewards may be used in making a navigational action selection at step 1625. In one example, it may be
30 determined that a reward identified at step 1621 (in response to at least one future action to a future state determined at step 1615 based on the second potential action determined at step 1607) may be higher than expected reward 1606. Based on this comparison, the second potential action determined at step 1607 may be selected rather than the first potential action determined at step 1605 despite expected reward 1606 being higher than expected reward 1608. In one example, the potential navigational action determined at step 1605 may include a merge in front of a detected target vehicle, while the potential
35 navigational action determined at step 1607 may include a merge behind the target vehicle. While the expected reward 1606 of merging in front of the target vehicle may be higher than the expected reward 1608 associated with merging behind the target vehicle, it may be determined that merging behind the target vehicle may result in a future state for which there may be action options yielding even higher

potential rewards than expected reward 1606, 1608, or other rewards based on available actions in response to a current, sensed navigational state.

[0371] Selection from among potential actions at step 1625 may be based on any suitable comparison of expected rewards (or any other metric or indicator of benefits associated with one potential action over another). In some cases, as described above, a second potential action may be selected over a first potential action if the second potential action is projected to provide at least one future action associated with an expected reward higher than a reward associated with the first potential action. In other cases, more complex comparisons may be employed. For example, rewards associated with action options in response to projected future states may be compared to more than one expected reward associated with a determined potential action.

[0372] In some scenarios, actions and expected rewards based on projected future states may affect selection of a potential action to a current state if at least one of the future actions is expected to yield a reward higher than any of the rewards expected as a result of the potential actions to a current state (e.g., expected rewards 1606, 1608, 1610, etc.). In some cases, the future action option that yields the highest expected reward (e.g., from among the expected rewards associated with potential actions to a sensed current state as well as from among expected rewards associated with potential future action options relative to potential future navigational states) may be used as a guide for selection of a potential action to a current navigational state. That is, after identifying a future action option yielding the highest expected reward (or a reward above a predetermined threshold, etc.), the potential action that would lead to the future state associated with the identified future action yielding the highest expected reward may be selected at step 1625.

[0373] In other cases, selection of available actions may be made based on determined differences between expected rewards. For example, a second potential action determined at step 1607 may be selected if a difference between an expected reward associated with a future action determined at step 1621 and expected reward 1606 is greater than a difference between expected reward 1608 and expected reward 1606 (assuming + sign differences). In another example, a second potential action determined at step 1607 may be selected if a difference between an expected reward associated with a future action determined at step 1621 and an expected reward associated with a future action determined at step 1619 is greater than a difference between expected reward 1608 and expected reward 1606.

[0374] Several examples have been described for selecting from among potential actions to a current navigational state. Any other suitable comparison technique or criteria, however, may be used for selecting an available action through long range planning based on action and reward analysis extending to projected future states. Additionally, while Fig. 16 represents two layers in the long range planning analysis (e.g., a first layer considering the rewards resulting from potential actions to a current state, and a second layer considering the rewards resulting from future action options in response to projected future states), analysis based on more layers may be possible. For example, rather than basing the long range planning analysis upon one or two layers, three, four or more layers of analysis could be used in selecting from among available potential actions in response to a current navigational state.

[0375] After a selection is made from among potential actions in response to a sensed navigational state, at step 1627, the at least one processor may cause at least one adjustment of a navigational actuator of the host vehicle in response to the selected potential navigational action. The navigational actuator may include any suitable device for controlling at least one aspect of the host vehicle. For example, the navigational actuator may include at least one of a steering mechanism, a
5 brake, or an accelerator.

[0376] Navigation Based on Inferred Aggression of Others

[0377] Target vehicles may be monitored through analysis of an acquired image stream to determine indicators of driving aggression. Aggression is described herein as a qualitative or
10 quantitative parameter, but other characteristics may be used: perceived level of attention (potential impairment of driver, distracted—cell phone, asleep, etc.). In some cases, a target vehicle may be deemed to have a defensive posture, and in some cases, the target vehicle may be determined to have a more aggressive posture. Navigational actions may be selected or developed based on indicators of aggression. For example, in some cases, the relative velocity, relative acceleration, increases in relative acceleration,
15 following distance, etc., relative to a host vehicle may be tracked to determine if the target vehicle is aggressive or defensive. If the target vehicle is determined to have a level of aggression exceeding a threshold, for example, the host vehicle may be inclined to give way to the target vehicle. A level of aggression of the target vehicle may also be discerned based on a determined behavior of the target vehicle relative to one or more obstacles in a path of or in a vicinity of the target vehicle (e.g., a leading
20 vehicle, obstacle in the road, traffic light, etc.).

[0378] As an introduction to this concept, an example experiment will be described with respect to a merger of the host vehicle into a roundabout, in which a navigational goal is to pass through and out of the roundabout. The situation may begin with the host vehicle approaches an entrance of the roundabout and may end with the host vehicle reaches an exit of the roundabout (e.g., the second exit).
25 Success may be measured based on whether the host vehicle maintains a safe distance from all other vehicles at all times, whether the host vehicle finishes the route as quickly as possible, and whether the host vehicle adheres to a smooth acceleration policy. In this illustration, N_T target vehicles may be placed at random on the roundabout. To model a blend of adversarial and typical behavior, with probability p , a target vehicle may be modeled by an “aggressive” driving policy, such that the aggressive target vehicle
30 accelerates when the host vehicle attempts to merge in front of the target vehicle. With probability $1-p$, the target vehicle may be modeled by a “defensive” driving policy, such that the target vehicle decelerates and lets the host vehicle merge in. In this experiment, $p = 0.5$, and the navigation system of the host vehicle may be provided with no information about the type of the other drivers. The types of other drivers may be chosen at random at the beginning of the episode.

[0379] The navigational state may be represented as the velocity and location of the host vehicle (the agent), and the locations, velocities, and accelerations of the target vehicles. Maintaining target acceleration observations may be important in order to differentiate between aggressive and defensive drivers based on the current state. All target vehicles may move on a one-dimensional curve

that outlines the roundabout path. The host vehicle may move on its own one-dimensional curve, which intersects the target vehicles' curve at the merging point, and this point is the origin of both curves. To model reasonable driving, the absolute value of all vehicles' accelerations may be upper bounded by a constant. Velocities may also be passed through a ReLU because driving backward is not allowed. Note, that by not allowing driving backwards, long-term planning may become a necessity, as the agent cannot regret on its past actions.

[0380] As described above, the next state, s_{t+1} , may be decomposed into a sum of a predictable part, $\hat{N}(s_t, a_t)$, and a non-predictable part, ν_t . The expression, $\hat{N}(s_t, a_t)$, may represent the dynamics of vehicle locations and velocities (which may be well-defined in a differentiable manner), while ν_t may represent the target vehicles' acceleration. It may be verified that $\hat{N}(s_t, a_t)$ can be expressed as a combination of ReLU functions over an affine transformation, hence it is differentiable with respect to s_t and a_t . The vector ν_t may be defined by a simulator in a non-differentiable manner, and may implement aggressive behavior for some targets and defensive behavior for other targets. Two frames from such a simulator are shown in Figs. 17A and 17B. In this example experiment, a host vehicle 1701 learned to slowdown as it approached the entrance of the roundabout. It also learned to give way to aggressive vehicles (e.g., vehicles 1703 and 1705), and to safely continue when merging in front of defensive vehicles (e.g., vehicles 1706, 1708, and 1710). In the example represented by Figs. 17A and 17B, the navigation system of host vehicle 1701 is not provided with the type of target vehicles. Rather, whether a particular vehicle is determined to be aggressive or defensive is determined through inference based on observed position and acceleration, for example, of the target vehicles. In Fig. 17A, based on position, velocity, and/or relative acceleration, host vehicle 1701 may determine that vehicle 1703 has an aggressive tendency and, therefore, host vehicle 1701 may stop and wait for target vehicle 1703 to pass rather than attempting to merge in front of target vehicle 1703. In Fig. 17B, however, target vehicle 1701 recognized that the target vehicle 1710 traveling behind vehicle 1703 exhibited defensive tendencies (again, based on observed position, velocity, and/or relative acceleration of vehicle 1710) and, therefore, completed a successful merge in front of target vehicle 1710 and behind target vehicle 1703.

[0381] Fig. 18 provides a flowchart representing an example algorithm for navigating a host vehicle based on predicted aggression of other vehicles. In the example of Fig. 18, a level of aggression associated with at least one target vehicle may be inferred based on observed behavior of the target vehicle relative to an object in the environment of the target vehicle. For example, at step 1801, at least one processing device (e.g., processing device 110) of the host vehicle navigation system may receive, from a camera associated with the host vehicle, a plurality of images representative of an environment of the host vehicle. At step 1803, analysis of one or more of the received images may enable the at least one processor to identify a target vehicle (e.g., vehicle 1703) in the environment of the host vehicle 1701. At step 1805, analysis of one or more of the received images may enable the at least one processing device to identify in the environment of the host vehicle at least one obstacle to the target vehicle. The object may include debris in a roadway, a stoplight/traffic light, a pedestrian, another vehicle (e.g., a vehicle traveling ahead of the target vehicle, a parked vehicle, etc.), a box in the roadway, a road barrier, a curb, or any

other type of object that may be encountered in an environment of the host vehicle. At step 1807, analysis of one or more of the received images may enable the at least one processing device to determine at least one navigational characteristic of the target vehicle relative to the at least one identified obstacle to the target vehicle.

5 [0382] Various navigational characteristics may be used to infer a level of aggression of a detected target vehicle in order to develop an appropriate navigational response to the target vehicle. For example, such navigational characteristics may include a relative acceleration between the target vehicle and the at least one identified obstacle, a distance of the target vehicle from the obstacle (e.g., a following distance of the target vehicle behind another vehicle), and/or a relative velocity between the target vehicle
10 and the obstacle, etc.

[0383] In some embodiments, the navigational characteristics of the target vehicles may be determined based on outputs from sensors associated with the host vehicle (e.g., radar, speed sensors, GPS, etc.). In some cases, however, the navigational characteristics of the target vehicles may be determined partially or fully based on analysis of images of an environment of the host vehicle. For
15 example, image analysis techniques described above and in, for example, U.S. Patent No. 9,168,868, which is incorporated herein by reference, may be used to recognize target vehicles within an environment of the host vehicle. And, monitoring a location of a target vehicle in the captured images over time and/or monitoring locations in the captured images of one or more features associated with the target vehicle (e.g., tail lights, head lights, bumper, wheels, etc.) may enable a determination of relative
20 distances, velocities, and/or accelerations between the target vehicles and the host vehicle or between the target vehicles and one or more other objects in an environment of the host vehicle.

[0384] An aggression level of an identified target vehicle may be inferred from any suitable observed navigational characteristic of the target vehicle or any combination of observed navigational characteristics. For example, a determination of aggressiveness may be made based on any observed
25 characteristic and one or more predetermined threshold levels or any other suitable qualitative or quantitative analysis. In some embodiments, a target vehicle may be deemed as aggressive if the target vehicle is observed to be following the host vehicle or another vehicle at a distance less than a predetermined aggressive distance threshold. On the other hand, a target vehicle observed to be following the host vehicle or another vehicle at a distance greater than a predetermined defensive distance threshold
30 may be deemed defensive. The predetermined aggressive distance threshold need not be the same as the predetermined defensive distance threshold. Additionally, either or both of the predetermined aggressive distance threshold and the predetermined defensive distance threshold may include a range of values, rather than a bright line value. Further, neither of the predetermined aggressive distance threshold nor the predetermined defensive distance threshold must be fixed. Rather these values, or ranges of values, may
35 shift over time, and different thresholds/ranges of threshold values may be applied based on observed characteristics of a target vehicle. For example, the thresholds applied may depend on one or more other characteristics of the target vehicle. Higher observed relative velocities and/or accelerations may warrant application of larger threshold values/ranges. Conversely, lower relative velocities and/or accelerations,

including zero relative velocities and/or accelerations, may warrant application of smaller distance threshold values/ranges in making the aggressive/defensive inference.

[0385] The aggressive/defensive inference may also be based on relative velocity and/or relative acceleration thresholds. A target vehicle may be deemed aggressive if its observed relative velocity and/or its relative acceleration with respect to another vehicle exceeds a predetermined level or range. A target vehicle may be deemed defensive if its observed relative velocity and/or its relative acceleration with respect to another vehicle falls below a predetermined level or range.

[0386] While the aggressive/defensive determination may be made based on any observed navigational characteristic alone, the determination may also depend on any combination of observed characteristics. For example, as noted above, in some cases, a target vehicle may be deemed aggressive based solely on an observation that it is following another vehicle at a distance below a certain threshold or range. In other cases, however, the target vehicle may be deemed aggressive if it both follows another vehicle at less than a predetermined amount (which may be the same as or different than the threshold applied where the determination is based on distance alone) and has a relative velocity and/or a relative acceleration of greater than a predetermined amount or range. Similarly, a target vehicle may be deemed defensive based solely on an observation that it is following another vehicle at a distance greater than a certain threshold or range. In other cases, however, the target vehicle may be deemed defensive if it both follows another vehicle at greater than a predetermined amount (which may be the same as or different than the threshold applied where the determination is based on distance alone) and has a relative velocity and/or a relative acceleration of less than a predetermined amount or range. System 100 may make an aggressive/defensive if, for example, a vehicle exceeds 0.5G acceleration or deceleration (e.g., jerk 5m/s³), a vehicle has a lateral acceleration of 0.5G in a lane change or on a curve, a vehicle causes another vehicle to do any of the above, a vehicle changes lanes and causes another vehicle to give way by more than 0.3G deceleration or jerk of 3m/s³, and/or a vehicle changes two lanes without stopping.

[0387] It should be understood that references to a quantity exceeding a range may indicate that the quantity either exceeds all values associated with the range or falls within the range. Similarly, references to a quantity falling below a range may indicate that the quantity either falls below all values associated with the range or falls within the range. Additionally, while the examples described for making an aggressive/defensive inference are described with respect to distance, relative acceleration, and relative velocity, any other suitable quantities may be used. For example, a time to collision may calculation may be used or any indirect indicator of distance, acceleration, and/or velocity of the target vehicle. It should also be noted that while the examples above focus on target vehicles relative to other vehicles, the aggressive/defensive inference may be made by observing the navigational characteristics of a target vehicle relative to any other type of obstacle (e.g., a pedestrian, road barrier, traffic light, debris, etc.).

[0388] Returning to the example shown in Figs. 17A and 17B, as host vehicle 1701 approaches the roundabout, the navigation system, including its at least one processing device, may receive a stream of images from a camera associated with the host vehicle. Based on analysis of one or

more of the received images, any of target vehicles 1703, 1705, 1706, 1708, and 1710 may be identified. Further, the navigation system may analyze the navigational characteristics of one or more of the identified target vehicles. The navigation system may recognize that the gap between target vehicles 1703 and 1705 represents the first opportunity for a potential merge into the roundabout. The navigation system may analyze target vehicle 1703 to determine indicators of aggression associated with target vehicle 1703. If target vehicle 1703 is deemed aggressive, then the host vehicle navigation system may choose to give way to vehicle 1703 rather than merging in front of vehicle 1703. On the other hand, if target vehicle 1703 is deemed defensive, then the host vehicle navigation system may attempt to complete a merge action ahead of vehicle 1703.

10 [0389] As host vehicle 1701 approaches the roundabout, the at least one processing device of the navigation system may analyze the captured images to determine navigational characteristics associated with target vehicle 1703. For example, based on the images, it may be determined that vehicle 1703 is following vehicle 1705 at a distance that provides a sufficient gap for the host vehicle 1701 to safely enter. Indeed, it may be determined that vehicle 1703 is following vehicle 1705 by a distance that exceeds an aggressive distance threshold, and therefore, based on this information, the host vehicle navigation system may be inclined to identify target vehicle 1703 as defensive. In some situations, however, more than one navigational characteristic of a target vehicle may be analyzed in making the aggressive/defensive determination, as discussed above. Furthering the analysis, the host vehicle navigation system may determine that, while target vehicle 1703 is following at a non-aggressive distance behind target vehicle 1705, vehicle 1703 has a relative velocity and/or a relative acceleration with respect to vehicle 1705 that exceeds one or more thresholds associated with aggressive behavior. Indeed, host vehicle 1701 may determine that target vehicle 1703 is accelerating relative to vehicle 1705 and closing the gap that exists between vehicles 1703 and 1705. Based on further analysis of the relative velocity, acceleration, and distance (and even a rate that the gap between vehicles 1703 and 1705 is closing), host vehicle 1701 may determine that target vehicle 1703 is behaving aggressively. Thus, while there may be a sufficient gap into which host vehicle may safely navigate, host vehicle 1701 may expect that a merge in front of target vehicle 1703 would result in an aggressively navigating vehicle directly behind the host vehicle. Further, target vehicle 1703 may be expected, based on the observed behavior through image analysis or other sensor output, that target vehicle 1703 would continue accelerating toward host vehicle 1701 or continuing toward host vehicle 1701 at a non-zero relative velocity if host vehicle 1701 was to merge in front of vehicle 1703. Such a situation may be undesirable from a safety perspective and may also result in discomfort to passengers of the host vehicle. For such reasons, host vehicle 1701 may choose to give way to vehicle 1703, as shown in Fig. 17B, and merge into the roundabout behind vehicle 1703 and in front of vehicle 1710, deemed defensive based on analysis of one or more of its navigational characteristics.

35 [0390] Returning to Fig. 18, at step 1809, the at least one processing device of the navigation system of the host vehicle may determine, based on the identified at least one navigational characteristic of the target vehicle relative to the identified obstacle, a navigational action for the host vehicle (e.g.,

merge in front of vehicle 1710 and behind vehicle 1703). To implement the navigational action (at step 1811), the at least one processing device may cause at least one adjustment of a navigational actuator of the host vehicle in response to the determined navigational action. For example, a brake may be applied in order to give way to vehicle 1703 in Fig. 17A, and an accelerator may be applied along with steering of the wheels of the host vehicle in order to cause the host vehicle to enter the roundabout behind vehicle 1703, as shown in Fig. 17B.

[0391] As described in the examples above, navigation of the host vehicle may be based on the navigational characteristics of a target vehicle relative to another vehicle or object. Additionally, navigation of the host vehicle may be based on navigational characteristics of the target vehicle alone without a particular reference to another vehicle or object. For example, at step 1807 of Fig. 18, analysis of a plurality of images captured from an environment of a host vehicle may enable determination of at least one navigational characteristic of an identified target vehicle indicative of a level of aggression associated with the target vehicle. The navigational characteristic may include a velocity, acceleration, etc. that need not be referenced with respect to another object or target vehicle in order to make an aggressive/defensive determination. For example, observed accelerations and/or velocities associated with a target vehicle that exceed a predetermined threshold or fall within or exceed a range of values may indicate aggressive behavior. Conversely, observed accelerations and/or velocities associated with a target vehicle that fall below a predetermined threshold or fall within or exceed a range of values may indicate defensive behavior.

[0392] Of course, in some instances the observed navigational characteristic (e.g., a location, distance, acceleration, etc.) may be referenced relative to the host vehicle in order to make the aggressive/defensive determination. For example, an observed navigational characteristic of the target vehicle indicative of a level of aggression associated with the target vehicle may include an increase in relative acceleration between the target vehicle and the host vehicle, a following distance of the target vehicle behind the host vehicle, a relative velocity between the target vehicle and the host vehicle, etc.

[0393] Navigation Based on Accident Liability Constraint

[0394] As described in the sections above, planned navigational actions may be tested against predetermined constraints to ensure compliance with certain rules. In some embodiments, this concept may be extended to considerations of potential accident liability. As discussed below, a primary goal of autonomous navigation is safety. As absolute safety may be impossible (e.g., at least because a particular host vehicle under autonomous control cannot control the other vehicles in its surroundings--it can only control its own actions), the use of potential accident liability as a consideration in autonomous navigation and, indeed, as a constraint to planned actions may help ensure that a particular autonomous vehicle does not take any actions that are deemed unsafe--e.g., those for which potential accident liability may attach to the host vehicle. If the host vehicle takes only actions that are safe and that are determined not to result in an accident of the host vehicle's own fault or responsibility, then desired levels of accident avoidance (e.g., fewer than 10^{-9} per hour of driving) may be achieved.

[0395] The challenges posed by most current approaches to autonomous driving include a lack of safety guarantees (or at least an inability to provide desired levels of safety), and also a lack of scalability. Consider the issue of guaranteeing multi-agent safe driving. As society will unlikely tolerate road accident fatalities caused by machines, an acceptable level of safety is paramount to the acceptance of autonomous vehicles. While a goal may be to provide zero accidents, this may be impossible since multiple agents are typically involved in an accident and one may envision situations where an accident occurs solely due to the blame of other agents. For example, as shown in Fig. 19, host vehicle 1901 drives on a multi-lane highway, and while host vehicle 1901 can control its own actions relative to the target vehicles 1903, 1905, 1907, and 1909, it cannot control the actions of the target vehicles surrounding it. As a result, host vehicle 1901 may be unable to avoid an accident with at least one of the target vehicles should vehicle 1905, for example, suddenly cut in to the host vehicle's lane on a collision course with the host vehicle. To address this difficulty, a typical response of autonomous vehicle practitioners is to resort to a statistical data-driven approach where safety validation becomes tighter as data over more mileage is collected.

[0396] To appreciate the problematic nature of a data-driven approach to safety, however, consider first that the probability of a fatality caused by an accident per one hour of (human) driving is known to be 10^{-6} . It is reasonable to assume that for society to accept machines to replace humans in the task of driving, the fatality rate should be reduced by three orders of magnitude, namely to a probability of 10^{-9} per hour. This estimate is similar to the assumed fatality rate of air bags and from aviation standards. For example, 10^{-9} is the probability that a wing will spontaneously detach from an aircraft in mid-air. Attempts to guarantee safety using a data-driven statistical approach that provides additional confidence with accumulating miles driven, however, is not practical. The amount of data required to guarantee a probability of 10^{-9} fatality per hour of driving is proportional to its inverse (i.e., 10^9 hours of data) which is roughly on the order of thirty billion miles. Moreover, a multi-agent system interacts with its environment and likely cannot be validated offline (unless a realistic simulator emulating real human driving with all its richness and complexities such as reckless driving is available--but the problem of validating the simulator would be even more difficult than creating a safe autonomous vehicle agent). And any change to the software of planning and control will require a new data collection of the same magnitude, which is clearly unwieldy and impractical. Further, developing a system through data invariably suffers from lack of interpretability and explainability of the actions being taken — if an autonomous vehicle (AV) has an accident resulting in a fatality, we need to know the reason. Consequently, a model-based approach to safety is required, but the existing “functional safety” and ASIL requirements in the automotive industry are not designed to cope with multi-agent environments.

[0397] A second primary challenge in developing a safe driving model for autonomous vehicles is the need for scalability. The premise underlying AV goes beyond “building a better world” and instead is based on the premise that mobility without a driver can be sustained at a lower cost than with a driver. This premise is invariably coupled with the notion of scalability — in the sense of supporting mass production of AVs (in the millions) and more importantly of supporting a negligible

incremental cost to enable driving in a new city. Therefore the cost of computing and sensing does matter, if AV is to be mass manufactured, the cost of validation and the ability to drive “everywhere” rather than in a select few cities is also a necessary requirement to sustain a business.

[0398] The issue with most current approaches lies in a “brute force” state of mind along three
5 axes: (i) the required “computing density,” (ii) the way high-definition maps are defined and created, and
(iii) the required specification from sensors. A brute-force approach goes against scalability and shifts the
weight towards a future in which unlimited on-board computing is ubiquitous, where the cost of building
and maintaining HD-maps becomes negligible and scalable, and exotic super advanced sensors would be
developed, productized to automotive grade, and at a negligible cost. A future for which any of the above
10 comes to fruition is indeed plausible but having all of the above hold is likely a low-probability event.
Thus, there is a need to provide a formal model that pieces together safety and scalability into an AV
program that society can accept and is scalable in the sense of supporting millions of cars driving
anywhere in the developed countries.

[0399] The disclosed embodiments represent a solution that may provide the target levels of
15 safety (or may even surpass safety targets) and may also be scalable to systems including millions of
autonomous vehicles (or more). On the safety front, a model called “Responsibility Sensitive Safety”
(RSS) is introduced that formalizes the notion of “accident blame,” is interpretable and explainable, and
incorporates a sense of “responsibility” into the actions of a robotic agent. The definition of RSS is
agnostic to the manner in which it is implemented — which is a key feature to facilitate a goal of creating
20 a convincing global safety model. RSS is motivated by the observation (as in Fig. 19) that agents play a
non-symmetrical role in an accident where typically only one of the agents is responsible for the accident
and therefore is to be responsible for it. The RSS model also includes a formal treatment of “cautious
driving” under limited sensing conditions where not all agents are always visible (due to occlusions, for
example). One primary goal of the RSS model is to guarantee that an agent will never make an accident
25 of its “blame” or for which it is responsible. A model may be useful only if it comes with an efficient
policy (e.g., a function that maps the “sensing state” to an action) that complies with RSS. For example,
an action that appears innocent at the current moment might lead to a catastrophic event in the far future
(“butterfly effect”). RSS may be useful for constructing a set of local constraints on the short-term future
that may guarantee (or at least virtually guarantee) that no accidents will happen in the future as a result
30 of the host vehicle’s actions.

[0400] Another contribution evolves around the introduction of a “semantic” language that
consists of units, measurements, and action space, and specification as to how they are incorporated into
planning, sensing and actuation of the AV. To get a sense of semantics, in this context, consider how a
human taking driving lessons is instructed to think about a “driving policy.” These instructions are not
35 geometric — they do not take the form “drive 13.7 meters at the current speed and then accelerate at a
rate of 0.8 m/s^2 ”. Instead, the instructions are of a semantic nature — “follow the car in front of you” or
“overtake that car on your left.” The typical language of human driving policy is about longitudinal and
lateral goals rather than through geometric units of acceleration vectors. A formal semantic language

may be useful on multiple fronts connected to the computational complexity of planning that do not scale up exponentially with time and number of agents, to the manner in which safety and comfort interact, to the way the computation of sensing is defined and the specification of sensor modalities and how they interact in a fusion methodology. A fusion methodology (based on the semantic language) may ensure that the RSS model achieves the required 10^{-9} probability of fatality, per one hour of driving, all while performing only offline validation over a dataset of the order of 10^5 hours of driving data.

[0401] For example, in a reinforcement learning setting, a Q function (e.g., a function evaluating the long term quality of performing an action $a \in A$ when the agent is at state $s \in S$; given such a Q-function, a natural choice of an action may be to pick the one with highest quality, $\pi(s) = \operatorname{argmax}_a Q(s, a)$) may be defined over a semantic space in which the number of trajectories to be inspected at any given time is bounded by 10^4 regardless of the time horizon used for planning. The signal to noise ratio in this space may be high, allowing for effective machine learning approaches to succeed in modeling the Q function. In the case of computation of sensing, semantics may allow for distinguishing between mistakes that affect safety versus those mistakes that affect the comfort of driving. We define a PAC model (Probably Approximate Correct (PAC)), borrowing Valiants PAC-learning terminology) for sensing which is tied to the Q-function and show how measurement mistakes are incorporated into planning in a manner that complies with RSS yet allows for optimization of the comfort of driving. The language of semantics may be important for the success of certain aspects of this model as other standard measures of error, such as error with respect to a global coordinate system, may not comply with the PAC sensing model. In addition, the semantic language may be an important enabler for defining HD-maps that can be constructed using low-bandwidth sensing data and thus be constructed through crowd-sourcing and support scalability.

[0402] To summarize, the disclosed embodiments may include a formal model that covers important ingredients of an AV: sense, plan, and act. The model may help ensure that from a planning perspective there will be no accident of the AV's own responsibility. And also through a PAC-sensing model, even with sensing errors, the described fusion methodology may require only offline data collection of a very reasonable magnitude to comply with the described safety model. Furthermore, the model may tie together safety and scalability through the language of semantics, thereby providing a complete methodology for a safe and scalable AV. Finally, it is worth noting that developing an accepted safety model that would be adopted by the industry and regulatory bodies may be a necessary condition for the success of AV.

[0403] The RSS model may generally follow a classic sense-plan-act robotic control methodology. The sensing system may be responsible for understanding a present state of the environment of a host vehicle. The planning part, which may be referred to as a "driving policy" and which may be implemented by a set of hard-coded instructions, through a trained system (e.g., a neural network), or a combination, may be responsible for determining what is the best next move in view of available options for accomplishing a driving goal (e.g., how to move from the left lane to a right lane in order to exit a highway). The acting portion is responsible for implementing the plan (e.g., the system of

actuators and one or more controllers for steering, accelerating, and/or braking, etc. a vehicle in order to implement a selected navigational action). The described embodiments below focus primarily on the sensing and planning parts.

[0404] Accidents may stem from sensing errors or planning errors. Planning is a multi-agent endeavor, as there are other road users (humans and machines) that react to actions of an AV. The described RSS model is designed to address safety for the planning part, among others. This may be referred to as multi-agent safety. In a statistical approach, estimation of the probability of planning errors may be done “online.” Namely, after every update of the software, billions of miles must be driven with the new version to provide an acceptable level of estimation of the frequency of planning errors. This is clearly infeasible. As an alternative, the RSS model may provide a 100% guarantee (or virtually 100% guarantee) that the planning module will not make mistakes of the AV’s blame (the notion of “blame” is formally defined). The RSS model may also provide an efficient means for its validation not reliant upon online testing.

[0405] Errors in a sensing system may be easier to validate, because sensing can be independent of the vehicle actions, and therefore we can validate the probability of a severe sensing error using “offline” data. But, even collecting offline data of more than 10^9 hours of driving is challenging. As part of the description of a disclosed sensing system, a fusion approach is described that can be validated using a significantly smaller amount of data.

[0406] The described RSS system may also be scalable to millions of cars. For example, the described semantic driving policy and applied safety constraints may be consistent with sensing and mapping requirements that can scale to millions of cars even in today’s technology.

[0407] for a foundational building block of such a system is a thorough safety definition, that is, a minimal standard to which AV systems may need to abide. In the following technical lemma, a statistical approach to validation of an AV system is shown to be infeasible, even for validating a simple claim such as “the system makes N accidents per hour”. This implies that a model-based safety definition is the only feasible tool for validating an AV system.

Lemma 1 *Let X be a probability space, and A be an event for which $\Pr(A) = p_1 < 0.1$. Assume we sample $m = \frac{1}{p_1}$ i.i.d. samples from X , and let $Z = \sum_{i=1}^m \mathbf{1}_{[x \in A]}$. Then*

$$\Pr(Z = 0) \geq e^{-2}.$$

Proof We use the inequality $1 - x \geq e^{-2x}$ (proven for completeness in Appendix A.1), to get

$$\Pr(Z = 0) = (1 - p_1)^m \geq e^{-2p_1 m} = e^{-2}.$$

Corollary 1 *Assume an AV system AV_1 makes an accident with small yet insufficient probability p_1 . Any deterministic validation procedure which is given $1/p_1$ samples, will, with constant probability, not distinguish between AV_1 and a different AV system AV_0 which never makes accidents.*

[0408] To gain perspective over the typical values for such probabilities, assume we desire an accident probability of 10^{-9} per hour, and a certain AV system provides only 10^{-8} probability. Even if the system obtains 10^8 hours of driving, there is constant probability that the validation process will not be able to indicate that the system is dangerous.

[0409] Finally, note that this difficulty is for invalidating a single, specific, dangerous AV system. A full *solution* cannot be viewed as a single system, as new versions, bug fixes, and updates will be necessary. Each change, even of a single line of code, generates a *new system* from a validator's perspective. Thus, a solution which is validated statistically, must do so online, over new samples after every small fix or change, to account for the shift in the distribution of states observed and arrived-at by the new system. Repeatedly and systematically obtaining such a huge number of samples (and even then, with constant probability, failing to validate the system), is infeasible.

[0410] Further, any statistical claim must be formalized to be measured. Claiming a statistical property over the number of accidents a system makes is significantly weaker than claiming "it drives in a safe manner." In order to say that, one must formally define what is safety.

[0411] Absolute Safety is Impossible

[0412] An action a taken by a car c may be deemed *absolutely safe* if no accident can follow the action at some future time. It can be seen that it is impossible to achieve absolute safety, by observing simple driving scenarios, for example, as depicted in Figure 19. From the perspective of vehicle 1901, no action can ensure that none of the surrounding cars will crash into it. Solving this problem by forbidding the autonomous car from being in such situations is also impossible. As every highway with more than two lanes will lead to it at some point, forbidding this scenario amounts to a requirement to remain in the garage. The implications might seem, at first glance, disappointing. Nothing is absolutely safe.

However, such a requirement for absolute safety, as defined above, may be too harsh, as evident by the fact that human drivers do not adhere to a requirement for absolute safety. Instead, humans behave according to a safety notion that depends on responsibility.

[0413] Responsibility-Sensitive Safety

[0414] An important aspect missing from the absolute safety concept is the non-symmetry of most accidents—it is usually one of the drivers who is responsible for a crash, and is to be blamed. In the example of Fig. 19, the central car 1901 is not to be blamed if the left car 1909, for example, suddenly drives into it. To formalize the fact that considering its lack of responsibility, a behavior of AV 1901 staying in its own lane can be considered safe. To do so, a formal concept of "accident blame" or accident responsibility, which can serve as the premise for a safe driving approach, is described.

[0415] As an example, consider the simple case of two cars c_f , c_r , driving at the same speed, one behind the other, along a straight road. Assume c_f the car at the front, suddenly brakes because of an obstacle appearing on the road, and manages to avoid it. Unfortunately, c_r did not keep enough of a distance from c_f is not able to respond in time, and crashes into c_f 's rear side. It is clear that the blame is on c_r ; it is the responsibility of the rear car to keep safe distance from the front car, and to be ready for unexpected, yet reasonable, braking.

[0416] Next, consider a much broader family of scenarios: driving in a multi-lane road, where cars can freely change lanes, cut into other cars' paths, drive at different speeds, and so on. To simplify the following discussion, assume a straight road on a planar surface, where the lateral, longitudinal axes are the x , y axes, respectively. This can be achieved, under mild conditions, by defining a homomorphism

between the actual curved road and a straight road. Additionally, consider a discrete time space. Definitions may aid in distinction between two intuitively different sets of cases: simple ones, where no significant lateral maneuver is performed, and more complex ones, involving lateral movement.

[0417] **Definition 1 (Car Corridor)** The corridor of a car c is the range $[c_{x,left}, c_{x,right}] \times [\pm\infty]$, where $c_{x,left}, c_{x,right}$ are the positions of the leftmost, rightmost corners of c .

[0418] **Definition 2 (Cut-in)** A car c_1 (car 2003 in Figs. 20A and 20B) cuts-in to car c_0 's (car 2001 in Figs. 20A and 20B) corridor at time t if it did not intersect c_0 's corridor at time $t - 1$, and does intersect it at time t .

[0419] A further distinction may be made between front/back parts of the corridor. The term "the direction of a cut-in" may describe movement in the direction of the relevant corridor boundary. These definitions may define cases with lateral movement. For the simple case where there is no such occurrence, such as the simple case of a car following another, the safe longitudinal distance is defined:

[0420] **Definition 3 (Safe longitudinal distance)** A longitudinal distance d (Fig. 21) between a car c_r (car 2103) and another car c_f (car 2105) that is in c_r 's frontal corridor is safe w.r.t. a response time p if for any braking command a , $|a| < a_{\max,brake}$, performed by c_f , if c_r will apply its maximal brake from time p until a full stop then it won't collide with c_f .

[0421] Lemma 2 below calculates d as a function of the velocities of c_r, c_f , the response time ρ , and the maximal acceleration $a_{\max,brake}$. Both ρ and $a_{\max,brake}$ are constants, which should be determined to some reasonable values by regulation.

[0422] **Lemma 2** Let c_r be a vehicle which is behind c_f on the longitudinal axis. Let $a_{\max,brake}, a_{\max,accel}$ be the maximal braking and acceleration commands, and let p be c_r 's response time. Let v_r, v_f be the longitudinal velocities of the cars, and let l_r, l_f be their lengths. Define $v_{p,max} = v_r + \rho \cdot a_{\max,accel}$, and define $T_r = p + \frac{v_{p,max}}{a_{\max,brake}}$ and $T_f = \frac{v_f}{a_{\max,brake}}$. Let $L = (l_r + l_f)/2$. Then, the minimal safe longitudinal distance for c_r is:

$$d_{\min} = \begin{cases} L & \text{if } T_r \leq T_f \\ L + T_f \left[(v_{p,max} - v_f) + \rho a_{\max,brake} \right] - \frac{\rho^2 a_{\max,brake}}{2} + \frac{(T_r - T_f)(v_{p,max} - (T_f - \rho) a_{\max,brake})}{2} & \text{otherwise} \end{cases}$$

[0423] **Proof** Let d_t be the distance at time t . To prevent an accident, we must have that $d_t > L$ for every t . To construct d_{\min} we need to find the tightest needed lower bound on d_0 . Clearly, d_0 must be at least L . As long as the two cars didn't stop after $T \geq \rho$ seconds, the velocity of the preceding car will be $v_f - T a_{\max,brake}$ while c_r 's velocity will be upper bounded by $v_{p,max} - (T - \rho) a_{\max,accel}$. So, the distance between the cars after T seconds will be lower bounded by:

$$\begin{aligned} d_T &:= d_0 + \frac{T}{2} (2v_f - T a_{\max,brake}) - \left[\rho v_{p,max} + \frac{T - \rho}{2} (2v_{p,max} - (T - \rho) a_{\max,brake}) \right] \\ &= d_0 + T \left[(v_f - v_{p,max}) - \rho a_{\max,brake} \right] + \frac{\rho^2 a_{\max,brake}}{2} \end{aligned}$$

[0424] Note that T_r is the time on which c_r arrives to a full stop (a velocity of 0) and T_f is the time on which the other vehicle arrives to a full stop. Note that $a_{\max, \text{brake}}(T_r - T_f) = v_{\rho, \max} - v_f + \rho a_{\max, \text{brake}}$, so if $T_r \leq T_f$ it suffices to require that $d_0 > L$. If $T_r > T_f$ then

$$d_{T_r} = d_0 + T_f [(v_f - v_{\rho, \max}) - \rho a_{\max, \text{brake}}] + \frac{\rho^2 a_{\max, \text{brake}}}{2} - \frac{(T_r - T_f)(v_{\rho, \max} - (T_f - \rho) a_{\max, \text{brake}})}{2}.$$

5 Requiring $d_{T_r} > L$ and rearranging terms concludes the proof.

[0425] Finally, a comparison operator is defined which allows comparisons with some notion of “margin”: when comparing lengths, velocities and so on, it is necessary to accept very similar quantities as “equal”.

Definition 4 (μ -comparison) The μ -comparison of two numbers a, b is $a >_{\mu} b$ if $a > b + \mu$, $a <_{\mu} b$ if $a < b - \mu$ and $a =_{\mu} b$ if $|a - b| \leq \mu$.

[0426] The comparisons (argmin, argmax, etc.) below are μ -comparisons for some suitable μ s. Assume an accident occurred between cars c_1, c_2 . To consider who is to blame for the accident, the relevant moment which needs to be examined is defined. This is some point in time which preceded the accident, and intuitively, was the “point of no return”; after it, nothing could be done to prevent the

15 accident.

[0427] **Definition 5 (Blame Time)** The Blame Time of an accident is the earliest time preceding the accident in which:

- there was an intersection between one of the cars and the other's corridor, and
- the longitudinal distance was not safe.

20 [0428] Clearly there is such a time, since at the moment of accident, both conditions hold.

Blame Times may be split into two separate categories:

- Ones in which a cut-in also occurs, namely, they are the first moment of intersection of one car and the other's corridor, and it's in a non-safe distance.
 - Ones in which a cut-in does not occur, namely, there was intersection with
- 25 the corridor already, in a safe longitudinal distance, and the distance had changed to unsafe at the Blame Time.

[0429] **Definition 6 (μ -Losing by Lateral Velocity)** Assume a cut-in occurs between cars c_1, c_2 . We say that c_1 μ -Loses by Lateral Velocity in case its lateral velocity w.r.t. the direction of the cut-in is higher by μ than that of c_2 .

30 [0430] It should be noted that the direction of the velocity is important: For example, velocities of $-1, 1$ (both cars crashing into each other) is a tie, however if the velocities are $1, 1 + \mu/2$, the one with positive direction towards the other car is to be blamed. Intuitively, this definition will allow us to blame a car which drives laterally very fast into another.

[0431] **Definition 7 ((μ_1, μ_2) -Winning by Lateral Position)** Assume a cut-in occurs between

35 cars c_1, c_2 . We say that c_1 (μ_1, μ_2) -Wins by Lateral Position in case its lateral position w.r.t. the cut-in lane's center (the center closest to the cut-in relevant corridor) is smaller than μ_1 (in absolute value), and smaller by μ_2 than that of c_2 .

[0432] Intuitively, we will not blame a car if it's very close to the lane center (μ_1), and much closer than the other car (by μ_2).

[0433] **Definition 8 (Blame)** *The Blame or responsibility for an accident between cars c_1 , c_2 , is a function of the state at the Blame Time, and is defined as follows:*

- 5
- *If the Blame Time is not a cut-in time, the blame is on the rear car.*
 - *If the Blame Time is also a cut-in time, the blame is on both cars, unless for one of the cars, w.l.o.g. c_1 , the two following conditions hold, for some predefined μ s:*
 - *It doesn't lose by Lateral Velocity,*
 - *It wins by Lateral Position.*

10 [0434] In that case, c_1 is spared. In other words, if an unsafe cut-in occurs, both cars are to blame, unless one of the cars is not (significantly) laterally faster, and is (significantly) closer to the lane center. By this, the desired behavior is captured: if following a car, keep a safe distance, and if cutting into a corridor of a car which simply drives in its own lane, do it only at a safe distance. An automated controller-based system for following the safety guidelines described above should not lead to overly
15 defensive driving, as discussed further below.

[0435] Dealing with Limited Sensing

[0436] After considering the highway example, a second example next addresses a problem of limited sensing. A very common human response, when blamed for an accident, falls into the "but I couldn't see him" category. It is, many times, true. Human sensing capabilities are limited, sometimes
20 because of an unaware decision to focus on a different part of the road, sometimes because of carelessness, and sometimes because of physical limitations - it is impossible to see a pedestrian hidden behind a parked car. Of those human limitations, advanced automatic sensing systems may only be subject to the latter: 360° view of the road, along with the fact that computers are never careless, puts them above human sensing capabilities. Returning to the "but I couldn't see him" example, a fitting
25 answer is "well, you should've been more careful." To formalize what is being careful with respect to limited sensing, consider the scenario, depicted in Fig. 22. Car 2201 (c_0) is trying to exit a parking lot, merging into a (possibly) busy road, but cannot see whether there are cars in the street because its view is obscured by building 2203. Assume that this is an urban, narrow street, with a speed limit of 30 km/h. A human driver's behavior is to slowly merge onto the road, obtaining more and more field of view, until
30 sensing limitations are eliminated. A significant moment in time should be defined - the first time the occluded object is exposed to us; after its exposure, one deal with it just like any other object that one can sense.

[0437] **Definition 9 (Exposure Time)** *The Exposure Time of an object is the first time in which we see it.*

35 [0438] **Definition 10 (Blame due to Unreasonable Speed)** *Assume that at the exposure time or after it, car c_1 (car 2205) was driving at speed $v > v_{limit}$, and c_0 wasn't doing so. Then, the blame is only on c_1 . We say that c_1 is blamed due to unreasonable speed.*

[0439] This extension allows c_0 to exit the parking lot safely. Using our previous responsibility-sensitive safety definitions, along with a dynamic v_{limit} definition (which uses the road conditions and speed limit, plus reasonable margins), the only necessity is to check whether in the worst case, as illustrated in the figure, the cut-in is in a safe longitudinal distance, while assuming that c_1 will not exceed v_{limit} . Intuitively, this encourages c_0 to drive slower and further from the occluder, thus slowly increasing its field of view and later allowing for safe merging into the street.

[0440] Having extended the accident responsibility definition to this basic case of limited sensing, a family of extensions may address similar cases. Simple assumptions as to what can be occluded (a potentially fast car cannot be occluded between two closely parked cars, but a pedestrian can), and what is the worst case maneuver it can perform (a pedestrian's v_{limit} is much smaller than that of a car), imply restrictions on driving - one must be prepared for the worst, and have the ability to respond if suddenly, the exposure time comes. A more elaborate example, in an urban scenario, can be taken from the scenario of a pedestrian which is possibly occluded by a parked car. Accident blame for accidents with a pedestrian may be defined:

[0441] **Definition 11 (Accident-with-Pedestrian Blame)** *The Accident-with-Pedestrian Blame is always on the car, unless one of the following three holds:*

- *the car hits the pedestrian with the car's side, and the lateral velocity of the car is smaller than μ , w.r.t. the direction of the hit.*
- *the pedestrian's velocity at the exposure time or later was larger than v_{limit} .*
- *the car is at complete stop.*

[0442] Informally, the car is not to blame only if a pedestrian runs into its side, while the car does not ride faster than μ into the pedestrian, or if the car is at stop, or if the pedestrian was running super-humanly fast, in some direction, not necessarily the hitting direction.

[0443] While the described system may not ensure absolute safety, it may lead to a scenario where very few (if any) accidents occur among autonomous vehicles. For example, where all cars (and other road users) are able to successfully verify that they will not be blamed for an accident as a result of an action taken, accidents may be eliminated. By definition, for every accident there is at least one responsible car. Thus, if no car takes an action for which it may be responsible for a resulting accident (according to the RSS model described above), there should never be any accidents, leading to the type of absolute safety or near absolute safety sought by unwieldy and impractical statistical methods.

[0444] Not all roads are of a simple structure. Some, like junctions and roundabouts, contain more complex situations, along with various right of way rules. Not all occluded objects are cars or pedestrians, with bicycles and motorcycles all legitimate road users to be considered. The principles introduced in this section may be extended to these additional cases.

[0445] Efficiently Validated Conditions for Responsibility-Sensitive Safety

[0446] This section discusses implementation aspects of RSS. To begin, it should be noted that an action that is performed now may have a butterfly effect that will lead to a chain of events with an accident after 10 minutes of driving, for example. A "brute-force" approach of checking all possible

future outcomes not only impractical, it is likely impossible. To overcome this challenge, the responsibility-sensitive safety definitions described above are now described together with computationally efficient methods to validate them.

[0447] Computationally Feasible Safety Verification.

5 [0448] The main mathematical tool for computationally feasible verification is “induction.” To prove a claim by induction, one begins with proving the claim for simple cases, and then, each induction step extends the proof to more and more involved cases. To illustrate how this induction tool can be helpful for safety verification, consider again the simple example of a car c_r following another, c_f (Fig. 21). The following constraint may be applied on the policy of c_r . At each time step t , the policy can
 10 pick any acceleration command such that even if c_f will apply a deceleration of $-a_{\max}$, the resulting distance between c_r and c_f at the next time step will be at least the safe longitudinal distance (defined in Definition 3 and Lemma 2). If no such action exists, c_r must apply the deceleration $-a_{\max}$. The following lemma uses induction to prove that any policy that adheres to the above constraints will never make an accident with c_f .

15 [0449] **Lemma 3** *Under the assumptions given in Definition 3, if the policy of c_r adheres to the constraints given above it will never make an accident with c_f .*

[0450] **Proof** The proof is by induction. For the induction base, start with an initial state in which the distance between the two cars is safe (according to Lemma 2). The induction step is as follows. Consider the distance between c_r and c_f at some time t . If there is an action that results in a safe distance
 20 (even with c_f making maximal deceleration), we are fine. If all actions cannot guarantee safe distance, let $t' < t$ be the maximal time in which we took an action which was not maximal deceleration. By the induction hypothesis, at time $t' + 1$, we were in a safe distance, and from there on we performed maximal deceleration. Hence, by the definition of safe distance, there was no crash from time t' till now, which concludes the proof.

25 [0451] The above example demonstrates a more general idea: there’s some emergency maneuver which can be performed by c_r in an extreme case, and lead it back to a “safe state.” It should be noted that the constraints on the policy we have described above depend on just one future time step, hence it can be verified in a computationally efficient manner.

[0452] To generalize those ideas of sufficient local properties for RSS, we firstly define a
 30 Default Emergency Policy (DEP), and use it as a building block for defining a local property of action-taking which we call “cautious”. It is then shown that taking only cautious commands is sufficient for RSS.

[0453] **Definition 12 (Default Emergency Policy)** *The Default Emergency Policy (DEP) is to apply maximum braking power, and maximum heading change towards 0 heading w.r.t. the lane. The
 35 maximum braking power and heading change are derived from physical parameters of the car (and maybe also from weather and road conditions). **Definition 13 (Safe state)** A state s is safe if performing DEP starting from it will not lead to an accident of our blame. As in the simple case of a car following another, we define a command to be cautious if it leads to a safe state. **Definition 14 (Cautious***

command) Suppose we are currently at state s_0 . A command a is cautious if the next state, s_1 , will be safe with respect to a set A of possible commands that other vehicles might perform now. The above definition depends on the worst-case commands, in the set A , other vehicles might perform. We will construct the set A based on reasonable upper bounds on maximum braking/acceleration and lateral movements.

5 [0454] The following theorem proves, by induction again, that if we only issue cautious commands then there will be no accidents of our blame. **Theorem 1** Assume that in time 0, c is in a safe state, and for every time step, c only issues cautious commands, where if no cautious command exists at some time step, c applies DEP. Then, c will never make accidents of its blame. **Proof** By induction. The base of the induction follows from the definition of a safe state and step from the definition of a cautious
10 command.

[0455] One benefit of this approach is that there may be no need to check infinite future, as we can quickly return to a safe state, and continue safely from there. Moreover, given the fact we will plan again at $t + 1$, and hence be able to perform DEP then if necessary, we should only check the command we are giving at time t , and not a possible longer plan we might have in mind - we can change that plan at
15 $t + 1$. Now, incorporating a learning component in a system, when it is verified at run time by this transparent model, is made possible. Finally, this local verification implies full future RSS, which is our desired goal. An implementation obstacle is that the cautiousness definition involves all trajectories another agent can perform until t_{brake} , which, even for moderate t_{brake} , is a huge space. To tackle this, we next turn to develop an efficiently computable way to verify cautiousness, and hence RSS, in a scalable
20 manner.

[0456] Efficient cautiousness verification

[0457] A first observation is that a state is not safe if and only if there exists a specific vehicle, \tilde{c} , which can perform commands from the set A which lead to an accident of our blame while we execute the DEP. Therefore, in a scene with a single target car, denoted \tilde{c} , and in the general case, the procedure
25 may be executed sequentially, for each of the other vehicles in the scene.

[0458] When considering a single target car, an action a is not cautious if and only if there is a sequence of commands for \tilde{c} , denoted $\tilde{a}_1, \dots, \tilde{a}_{t_{\text{brake}}}$, all in the set A , that results in an accident of c 's blame. As already proven, if at time 0, it holds that \tilde{c} is in the frontal corridor of c , there is a simple way to check the cautiousness of a — we just need to verify that even if \tilde{c} will apply maximal brake for one time step
30 (and we'll perform a), the resulting longitudinal distance will remain safe. The lemma below gives a sufficient condition for cautiousness in the more involved cases, where lateral maneuvers are to be considered too.

[0459] **Lemma 4** Assume that at time $T = 0$, \tilde{c} is not in the frontal corridor of c . Then, if at every $T \in (0, t_{\text{brake}}]$, there is no non-safe cut in of c 's blame, then a is cautious.

35 [0460] **Proof** Suppose that a is not cautious, namely, there exists $\tilde{a}_1, \dots, \tilde{a}_{t_{\text{brake}}}$ that leads to an accident of c 's blame. Before the accident there must be a cut-in time, T . Assume first that $T > 0$. If this cut-in was at a safe longitudinal distance, then there cannot be an accident of our blame due to the fact that DEP is executed by deceleration of $-a_{\text{max}}$ and based on the definition of safe longitudinal distance

(and we assume here that the response time ρ is larger than the time resolution of steps). If the cut-in was non-safe, by the assumption of the lemma it was not of c 's blame, hence the accident is also not of c 's blame.

[0461] Finally, if $T < 0$, by the assumption of the lemma, at the moment of cutting, \tilde{c} was at the back corridor of c . By induction, c performed only safe cut-ins in the past, and hence either the cut-in was safe or it was c 's blame. In both cases, the current accident is not of c 's blame.

[0462] In light of Lemma 4, the problem of checking whether there can be a non-safe cut in of c 's blame remains. An efficient algorithm is presented below for checking the possibility of a non-safe cut-in at time t . To validate the entire trajectory we discretize the time interval $[0, t_{\text{brake}}]$ and apply the algorithm on all time intervals (with a slightly larger value of ρ in the definition of safe distance to ensure that the discretization doesn't hurt). Let \tilde{c}_{diag} be the length of a diagonal of a minimal rectangle bounding \tilde{c} . For each time $t \in [0, t_{\text{brake}}]$, define $c_{\text{length}}(t)$ to be the longitudinal "span" of c in time t , and let $L(t) = \frac{\tilde{c}_{\text{diag}} + c_{\text{length}}[t]}{2}$. Define $c_{\text{width}}[t]$ in similar fashion and $W(t) = \frac{\tilde{c}_{\text{diag}} + c_{\text{width}}[t]}{2}$.

Algorithm 1: Check the possibility of a non-safe cut-in at time t

input:

$\tilde{y}[0], \tilde{v}_y[0], \tilde{x}[0], \tilde{v}_x[0]$ are longitudinal/lateral position/velocity of \tilde{c} at time 0
 $y[t], v_y[t], x[t], v_x[t]$ are longitudinal/lateral position/velocity of c at time t
 $a_{y,\min}, a_{y,\max}$ are longitudinal acceleration boundaries
 $a_{x,\max}$ is a lateral acceleration absolute value boundary
 $L = L(t), W = W(t)$

check longitudinal feasibility:

let $\tilde{y}_{\min} = \tilde{y}[0] + \tilde{v}_y[0]t + \frac{1}{2}a_{y,\min}t^2$
 let $\tilde{y}_{\max} = \tilde{y}[0] + \tilde{v}_y[0]t + \frac{1}{2}a_{y,\max}t^2$
 if $[\tilde{y}_{\min}, \tilde{y}_{\max}] \cap [y[t] - L, y[t] + L] \neq \emptyset$
 continue to check lateral feasibility
 else if $\tilde{y}_{\min} > y[t] + L$ and $(\tilde{y}_{\min}, \tilde{v}_y[0] + a_{y,\min}t)$ is not longitudinally safe w.r.t. $(y[t], v_y[t])$
 continue to check lateral feasibility
 else if $\tilde{y}_{\max} < y[t] - L$ and $(\tilde{y}_{\max}, \tilde{v}_y[0] + a_{y,\max}t)$ is not longitudinally safe w.r.t. $(y[t], v_y[t])$
 continue to check lateral feasibility
 return "non-feasible"

check lateral feasibility:

w.l.o.g. assume $x[t] = 0$ and $\tilde{x}[0] \leq 0$, and $\tilde{v}_x[0] \geq 0$.
 if a position $x[t]$ is not considered (μ_1, μ_2) -Winning by Lateral Position w.r.t. a position $x[t] - W$
 w.l.o.g. assume $v_x[t] = -\mu$, as in Definition 6.
 else
 w.l.o.g. assume $v_x[t] = -2\mu$, as in Definition 6.
 let $t_{\text{top}} = 0.5(t - \tilde{v}_x[0]/a_{x,\max})$
 if $t_{\text{top}} < 0$
 return "non-feasible"
 let $x_{\max} = \tilde{x}[0] + 2(\tilde{v}[0]t_{\text{top}} + 0.5a_{x,\max}t_{\text{top}}^2) + \tilde{v}_x[0]^2/(2a_{x,\max})$
 if $x_{\max} < -W$
 return "non-feasible"
 return "feasible"

[0463] The following theorem proves the correctness of the above algorithm.

[0464] **Theorem 2** *If Algorithm 1 returns “non-feasible” then there cannot be a non-safe cut-in of the ego vehicle’s blame at time t .* To prove the theorem, we rely on the following key lemmas, which prove the correctness of the two building blocks of Algorithm 1. Start with the longitudinal feasibility:

5 [0465] **Lemma 5** *Under the notation of Algorithm 1, if the **check longitudinal feasibility** procedure is concluded by returning “non-feasible”, then there cannot be a non-safe cut-in of the ego vehicle’s blame at time t .*

[0466] **Proof** Ignoring the lateral aspect of a cut-in maneuver, we examine the mere possibility that the longitudinal distance between c and \tilde{c} will be unsafe. It is clear that the positions $(\tilde{y}_{\min}, \tilde{y}_{\max})$ are bounding the position which can be attained by \tilde{c} at time t . By the fact $[\tilde{y}_{\min}, \tilde{y}_{\max}] \cap [y[t] - L, y[t] + L] = \emptyset$, we obtain that any longitudinally non-safe distance which is attainable, is $\geq L$. Assume $\tilde{y}_{\min} > y[t] + L$, and assume by contradiction that a non-safe longitudinal position and velocity, denoted $\tilde{y}_{\text{bad}}[t]$, $\tilde{v}_{\text{ybad}}[t]$, are attainable using acceleration commands bounded by $a_{y,\min}, a_{y,\max}$. By definition of \tilde{y}_{\min} , we have $\tilde{y}_{\text{bad}}[t] > \tilde{y}_{\min}$, and hence the distance between the cars is larger, namely $\tilde{y}_{\text{bad}}[t] - (y[t] + L) > \tilde{y}_{\min} - (y[t] + L)$. Since $(\tilde{y}_{\min}, \tilde{v}_y[0] + a_{y,\min}t)$ is longitudinally safe w.r.t. $(y[t], v_y[t])$, by definition of longitudinal non-safety, it follows that the attained velocity $\tilde{v}_{\text{ybad}}[t]$ must be smaller than $\tilde{v}_y[0] + a_{y,\min}t$. However, it is clear that in order to achieve lesser speed, \tilde{c} must use average acceleration which is lesser than $a_{y,\min}$ throughout the time window $[0, t]$, thus contradicting the fact that the longitudinal non-safety was attained using commands bounded by $a_{y,\min}, a_{y,\max}$. By considering a symmetrical argument for the case $\tilde{y}_{\max} < y[t] - L$, the proof is completed.

10
15
20

[0467] Next, the lateral feasibility.

[0468] **Lemma 6** *Under the notation of Algorithm 1, if the **check lateral feasibility** procedure is concluded by returning “non-feasible”, then there cannot be a non-safe cut-in of the ego vehicle’s blame at time t .*

25 [0469] **Proof** First, it is clear that there is no loss of generality by the assumptions $x[t] = 0$, $\tilde{x}[0] \leq 0$ and the ones regarding $v_x[t]$, by a simple change of coordinates and consideration of relative velocity. Moreover, by similar arguments it is simple to extend to the case when $\tilde{v}_x[0] \leq 0$.

[0470] Note that the positions of the cars involved in the cut-in, which in our case, are $(x[t], x[t] - W)$, imply some $(\mu 1, \mu 2)$ -Winning by Lateral Position property, affecting the blame. By our assumptions over $v_x[t]$, we obtain that the maximal lateral velocity \tilde{c} may use at time t , under assumption that c will be blamed, is 0: either in order to μ -“tie” lateral velocity (in the case c does not $(\mu 1, \mu 2)$ -Win by Lateral Position, this can be enough in order to put the blame on it), or to μ -Win lateral velocity (in the case c does $(\mu 1, \mu 2)$ -Win by Lateral Position, this is necessary in order to put the blame on it). It is left to check whether exists a maneuver starting at $\tilde{v}_x[0]$, ending at $\tilde{v}_x[t] = 0$, using lateral accelerations bounded by $a_{x,\max}$, with final position $\tilde{x}[t] \geq x[t] - W$. In words, a cut-in which ends at the desired lateral velocity, 0.

30
35

[0471] Recall the definition $t_{\text{top}} = 0.5(t - \tilde{v}_x[0]/a_{x,\max})$ from the algorithm. Assume $t_{\text{top}} < 0$. This implies that the time needed by \tilde{c} in order to reach lateral velocity 0 when using maximal lateral

acceleration, namely, $\tilde{v}_x[0]/a_{x,\max}$, is lesser than t . This implies that there is no manoeuvre which it can perform in order to reach the desired velocity in time, and hence no problematic maneuver exists.

Therefore, if the procedure returned “non-feasible” due to $t_{\text{top}} < 0$, indeed, there is no feasibility of a non-safe cut-in of c 's blame.

5 [0472] Consider the case $t_{\text{top}} > 0$. Then, the procedure returned “non-feasible” due to $x_{\max} < -W$. Consider a family of lateral velocity profiles for \tilde{c} in the time range $[0, t]$, denoted $U = \{u_a : a > \tilde{v}_x[0]/t\}$ and parameterized by a . We define, for each a , in similar fashion to the one used in the algorithm, $t_{\text{top}}(a) := 0.5(t - \tilde{v}_x[0]/a)$. Note that $t_{\text{top}}(a) > 0$ for all $a > \tilde{v}_x[0]/t$. We now define the velocity profile u_a for all times $t' \in [0, t]$ as follows:

$$10 \quad u_a(t') = \begin{cases} \tilde{v}_x[0] + a \cdot t' & t' < t_{\text{top}}(a) \\ \tilde{v}_x[0] + a \cdot (2t_{\text{top}}(a) - t') & t' \geq t_{\text{top}}(a) \end{cases}$$

[0473] First, it can be seen that u_a satisfies the constraints $u_a(0) = \tilde{v}_x[0]$, $u_a(t) = \tilde{v}_x[t]$. Second, the distance travelled while using u_a can be calculated, as this amounts to integration of a piecewise linear function. Define the arrived-at position as \tilde{x}_{u_a} , and note that x_{\max} defined in the algorithm is precisely $x_{u_{a_{x,\max}}}$. Third, it can be seen that the travelled distance is monotonically increasing with a , and is

15 unbounded. Hence, for any desired final position $x > \tilde{x}_{u_{\tilde{v}_x[0]/t}}$, there exists a value of a for which $x = \tilde{x}_{u_a}$. In particular, for $x = x[t] - W$, such a value exists, and we denote it by a_{cut} .

[0474] Observe that since x_{\max} , defined in the algorithm, is $< x[t] - W$, we have that $a_{\text{cut}} > a_{x,\max}$. This is not sufficient in order to show no valid maneuver can lead to position $\geq x[t] - W$; this is implied only for members of the family U . We now prove that even outside of U , all velocity profiles which attain a final position of $x[t] - W$ must use an acceleration value of at least a_{cut} on the way, making them invalid, and hence completing the proof.

20 [0475] Assume some velocity profile u satisfies the boundary constraints $u(0) = \tilde{v}_x[0]$, $u(t) = \tilde{v}_x[t]$. Moreover, assume it attains some final position $\tilde{x}_u > x[t] - W$. We thus obtain:

$$\int_0^t u(\tau) d\tau \geq \int_0^t u_{a_{\text{cut}}}(\tau) d\tau.$$

25 [0476] Assume $u \geq u_{a_{\text{cut}}}$ for all τ . In particular, $u(t_{\text{top}}(a_{\text{cut}})) \geq u_{a_{\text{cut}}}(t_{\text{top}}(a_{\text{cut}}))$. From the mean value theorem, there exists $\zeta \in [0, t_{\text{top}}(a_{\text{cut}})]$ s.t.

$$u'(\zeta) = \frac{u(t_{\text{top}}(a_{\text{cut}})) - u(0)}{t_{\text{top}}(a_{\text{cut}})} = \frac{u(t_{\text{top}}(a_{\text{cut}})) - u_{a_{\text{cut}}}(0)}{t_{\text{top}}(a_{\text{cut}})} \geq \frac{u_{a_{\text{cut}}}(t_{\text{top}}(a_{\text{cut}})) - u_{a_{\text{cut}}}(0)}{t_{\text{top}}(a_{\text{cut}})} = a_{\text{cut}} >$$

$a_{x,\max}$,

implying infeasibility of u , as it uses acceleration (that is, u' , the derivative of the velocity) which exceeds

30 $a_{x,\max}$.

[0477] Now, assume $u \geq u_{a_{\text{cut}}}$ does not hold for all τ . Then, due to the fact

$\int_0^t u(\tau) d\tau \geq \int_0^t u_{a_{\text{cut}}}(\tau) d\tau$, there must be a point where $u > u_{a_{\text{cut}}}$. If such point τ_{large} exists in $[0, t_{\text{top}}(a_{\text{cut}})]$,

then we can easily use the mean value theorem in the same manner as above, to obtain $\zeta \in [0, \tau_{\text{large}}]$ where too large an acceleration is used. If such point only exists in $[t_{\text{top}}(a_{\text{cut}}), t]$, similar argument will give us a

point $\zeta_0 [\tau_{\text{large}}, t]$ in which an acceleration value lesser than $-\alpha_{x,\text{max}}$ was used, concluding the proof.

Equipped with the above lemmas, Theorem 2's proof is immediate.

[0478] Safety Verification - Occlusions

[0479] In similar fashion to dealing with observed objects, we can define the extension for
5 cautiousness w.r.t. occluded objects, with a similar theorem to Theorem 1, proving that cautiousness implies there are never accidents of our blame.

[0480] **Definition 15 (Cautiousness w.r.t. Occluded Objects)** *A command given at time t is Cautious w.r.t. Occluded Objects if in the case that the exposure time of the object is $t + 1$, and we command a Default Emergency Policy (DEP) at $t + 1$, there will not be an accident of our blame.*

10 [0481] **Lemma 7** *If we only give cautious, w.r.t. occluded objects and non occluded objects, commands, there will never be an accident of our blame.*

[0482] **Proof** Assume that an accident of our blame occurred at time t , with the exposure time being $t' \leq t$. By cautiousness assumption, the command given at $t' - 1$ allowed us to command a DEP at t' without being blamed for an accident. As there was an accident of our blame, we apparently did not
15 command a DEP at time t' . But from t' on, we were safe w.r.t. non occluded objects, hence the command we gave was safe, and there was no accident of our blame.

[0483] Here too, we provide efficient ways to checking cautiousness with respect to a worst-case assumption over occluded objects, allowing for feasible, scalable, RSS.

[0484] Driving Policy

20 [0485] A driving policy is a mapping from a sensing state (a description of the world around us) into a driving command (e.g., the command is lateral and longitudinal accelerations for the coming second, which determines where and at what speed should the car be in one second from now). The driving command is passed to a controller, which aims at actually moving the car to the desired position/speed.

25 [0486] In the previous sections a formal safety model and proposed constraints on the commands issued by the driving policy that guarantee safety were described. The constraints on safety are designed for extreme cases. Typically, we do not want to even need these constraints, and would like to construct a driving policy that leads to a comfortable ride. The focus of this section is on how to build an efficient driving policy, in particular, one that requires computational resources that can scale to
30 millions of cars. For now, this discussion does not address the issue of how to obtain the sensing state and assume an utopic sensing state, that faithfully represents the world around us without any limitations. Later sections discuss the effect of inaccuracies in the sensing state on the driving policy.

[0487] The problem of defining a driving policy is cast in the language of Reinforcement Learning (RL), as discussed in the sections above. At each iteration of RL, an agent observes a state
35 describing the world, denoted s_t , and should pick an action, denoted a_t , based on a policy function, π , that maps states into actions. As a result of its action and other factors out of its control (such as the actions of other agents), the state of the world is changed to s_{t+1} . We denote a (state, action) sequence by $\bar{s} = ((s_1, a_1), (s_2, a_2), \dots, (s_{\text{len}(\bar{s})}, a_{\text{len}(\bar{s})}))$. Every policy induces a probability function over (state, action)

sequences. This probability function may be affected by the actions taken by the agent, but also depends on the environment (and in particular, on how other agents behave). We denote by P_π the probability over (state, action) sequences induced by π . The quality of a policy is defined to be $\mathbb{E}_{\bar{s} \sim P_\pi}[\rho(\bar{s})]$, where $\rho(\bar{s})$ is a reward function that measures how good the sequence \bar{s} is. In most case, $\rho(\bar{s})$ takes the form $\rho(\bar{s}) =$

5 $\sum_{t=1}^{\text{len}(\bar{s})} \rho(s_t, a_t)$, where $\rho(s, a)$ is an instantaneous reward function that measures the immediate quality of being at state s and performing action a . For simplicity, we stick to this simpler case.

[0488] To cast the driving policy problem in the above RL language, let s_t be some representation of the road, and the positions, velocities, and accelerations, of the ego vehicle as well as other road users. Let a_t be a lateral and longitudinal acceleration command. The next state, s_{t+1} , depends
10 on a_t as well as on how the other agents will behave. The instantaneous reward, $\rho(s_t, a_t)$, may depend on the relative position/velocities/acceleration to other cars, the difference between our speed and the desired speed, whether we follow the desired route, whether our acceleration is comfortable etc.

[0489] One challenge in deciding what action should the policy take at time t stems from the fact that one needs to estimate the long term effect of this action on the reward. For example, in the
15 context of driving policy, an action that is taken at time t may seem a good action for the present (that is, the reward value $\rho(s_t, a_t)$ is good), but might lead to an accident after 5 seconds (that is, the reward value in 5 seconds would be catastrophic). We therefore need to estimate the long term quality of performing an action a when the agent is at state s . This is often called the Q -function, namely, $Q(s, a)$ should reflect the long term quality of performing action a at time s . Given such a Q -function, the natural choice of an
20 action is to pick the one with highest quality, $\pi(s) = \text{argmax}_a Q(s, a)$.

[0490] The immediate questions are how to define Q and how to evaluate Q efficiently. Let us first make the (completely non-realistic) simplifying assumption that s_{t+1} is some deterministic function of (s_t, a_t) , namely, $s_{t+1} = f(s_t, a_t)$. One familiar with Markov Decision Processes (MDPs), will notice that this assumption is even stronger than the Markovian assumption of MDPs (i.e., that s_{t+1} is conditionally
25 independent of the past given (s_t, a_t)). As noted in [5], even the Markovian assumption is not adequate for multi-agent scenarios, such as driving, and we will therefore later relax the assumption.

[0491] Under this simplifying assumption, given s_t , for every sequence of decisions for T steps, (a_t, \dots, a_{t+T}) , we can calculate exactly the future states $(s_{t+1}, \dots, s_{t+T+1})$ as well as the reward values for times t, \dots, T . Summarizing all these reward values into a single number, e.g. by taking their
30 sum $\sum_{\tau=t}^T \rho(s_\tau, a_\tau)$, we can define $Q(s, a)$ as follows:

$$Q(s, a) = \max_{(a_t, \dots, a_{t+T})} \sum_{\tau=t}^T \rho(s_\tau, a_\tau) \quad \text{s.t.} \quad s_t = s, a_t = a, \forall \tau, s_{\tau+1} = f(s_\tau, a_\tau)$$

[0492] That is, $Q(s, a)$ is the best future we can hope for, if we are currently at state s and immediately perform action a .

[0493] Let us discuss how Q may be calculated. The first idea is to discretize the set of
35 possible actions, A , into a finite set \hat{A} , and simply traverse all action sequences in the discretized set. Then, the runtime is dominated by the number of discrete action sequences, $|\hat{A}|^T$. If \hat{A} represents 10 lateral accelerations and 10 longitudinal accelerations, we obtain 100^T possibilities, which becomes infeasible

even for small values of T . While there are heuristics for speeding up the search (e.g. coarse-to-fine search), this brute-force approach requires tremendous computational power.

[0494] The parameter T is often called the “time horizon of planning”, and it controls a natural tradeoff between computation time and quality of evaluation — the larger T is, the better our evaluation of the current action (since we explicitly examine its effect deeper into the future), but on the other hand, a larger T increases the computation time exponentially. To understand why we may need a large value of T , consider a scenario in which we are 200 meters before a highway exit and we should take it. When the time horizon is long enough, the cumulative reward will indicate if at some time τ between t and $t + T$ we have arrived to the exit lane. On the other hand, for a short time horizon, even if we perform the right immediate action we will not know if it will lead us eventually to the exit lane.

[0495] A different approach attempts to perform offline calculations in order to construct an approximation of Q , denoted \hat{Q} , and then during the online run of the policy, use \hat{Q} as an approximation to Q , without explicitly rolling out the future. One way to construct such an approximation is to discretize both the action domain and the state domain. Denote by \hat{A} , \hat{S} these discretized sets. An offline calculation may evaluate the value of $Q(s, a)$ for every $(s, a) \in \hat{S} \times \hat{A}$. Then, for every $a \in \hat{A}$ we define $\hat{Q}(s, a)$ to be $Q(s, a)$ for $s = \operatorname{argmin}_{s_0 \in \hat{S}} \|s - s_0\|$. Furthermore, based on the pioneering work of Bellman [2, 3], we can calculate $Q(s, a)$ for every $(s, a) \in \hat{S} \times \hat{A}$, based on dynamic programming procedures (such as the Value Iteration algorithm), and under our assumptions, the total runtime is order of $T |\hat{A}| |\hat{S}|$. The main problem with this approach is that in any reasonable approximation, \hat{S} is extremely large (due to the curse of dimensionality). Indeed, the sensing state should represent 6 parameters for every other relevant vehicle in the sense — the longitudinal and lateral position, velocity, and acceleration. Even if we discretize each dimension to only 10 values (a very crude discretization), since we have 6 dimensions, to describe a single car we need 10^6 states, and to describe k cars we need 10^{6k} states. This leads to unrealistic memory requirements for storing the values of Q for every $(s, a) \in \hat{S} \times \hat{A}$.

[0496] One approach for dealing with this curse of dimensionality is to restrict Q to come from a restricted class of functions (often called a hypothesis class), such as linear functions over manually determined features or deep neural networks. For example, consider a deep neural network that approximates Q in the context of playing Atari games. This leads to a resource-efficient solution, provided that the class of functions that approximate Q can be evaluated efficiently. However, there are several disadvantages of this approach. First, it is not known if the chosen class of functions contains a good approximation to the desired Q function. Second, even if such function exists, it is not known if existing algorithms will manage to learn it efficiently. So far, there are not many success stories for learning a Q function for complicated multi-agent problems, such as the ones we are facing in driving. There are several theoretical reasons why this task is difficult. As mentioned regarding the Markovian assumption, underlying existing methods are problematic. But, a more severe problem is a very small signal-to-noise ratio due to the time resolution of decision making, as explained below.

[0497] Consider a simple scenario in which a vehicle needs to change lane in order to take a highway exit in 200 meters and the road is currently empty. The best decision is to start making the lane

change. Decisions may be made every 0.1 second, so at the current time t , the best value of $Q(s, a)$ should be for the action a corresponding to a small lateral acceleration to the right. Consider the action a' that corresponds to zero lateral acceleration. Since there is a very little difference between starting the change lane now, or in 0.1 seconds, the values of $Q(s, a)$ and $Q(s, a')$ are almost the same. In other words, there is very little *advantage* for picking a over a' . On the other hand, since we are using a function approximation for Q , and since there is noise in measuring the state s_t , it is likely that our approximation to the Q value is noisy. This yields a very small signal-to-noise ratio, which leads to an extremely slow learning, especially for stochastic learning algorithms which are heavily used for the neural networks approximation class. However, as noted in, this problem is not a property of any particular function approximation class, but rather, it is inherent in the definition of the Q function.

[0498] In summary, available approaches can be roughly divided into two camps. The first one is the brute-force approach which includes searching over many sequences of actions or discretizing the sensing state domain and maintaining a huge table in memory. This approach can lead to a very accurate approximation of Q but requires unlimited resources, either in terms of computation time or in terms of memory. The second one is a resource efficient approach in which we either search for short sequences of actions or we apply a function approximation to Q . In both cases, we pay by having a less accurate approximation of Q that might lead to poor decisions.

[0499] The approach described herein includes constructing a Q function that is both resource-efficient and accurate is to depart from geometrical actions and to adapt a semantic action space, as described in the next subsection.

[0500] *Semantic approach*

[0501] As a basis for the disclosed semantic approach, consider a teenager that just got his driving license. His father sits next to him and gives him “driving policy” instructions. These instructions are not geometric — they do not take the form “drive 13.7 meters at the current speed and then accelerate at a rate of 0.8 m/s^2 ”. Instead, the instructions are of semantic nature — “follow the car in front of you” or “quickly overtake that car on your left.” We formalize a semantic language for such instructions, and use them as a semantic action space. We then define the Q function over the semantic action space. We show that a semantic action can have a very long time horizon, which allows us to estimate $Q(s, a)$ without planning for many future semantic actions. Yes, the total number of semantic actions is still small. This allows us to obtain an accurate estimation of the Q function while still being resource efficient. Furthermore, as we show later, we combine learning techniques for further improving the quality function, while not suffering from a small signal-to-noise ratio due to a significant difference between different semantic actions.

[0502] Now define a semantic action space. The main idea is to define lateral and longitudinal goals, as well as the aggressiveness level of achieving them. Lateral goals are desired positions in lane coordinate system (e.g., “my goal is to be in the center of lane number 2”). Longitudinal goals are of three types. The first is relative position and speed with respect to other vehicles (e.g., “my goal is to be behind car number 3, at its same speed, and at a distance of 2 seconds from it”). The second is a speed target

(e.g., “drive at the allowed speed for this road times 110%”). The third is a speed constraint at a certain position (e.g., when approaching a junction, “speed of 0 at the stop line”, or when passing a sharp curve, “speed of at most 60 km/h at a certain position on the curve”). For the third option, we can instead apply a “speed profile” (few discrete points on the route and the desired speed at each of them). A reasonable
 5 number of lateral goals is bounded by $16 = 4 \times 4$ (4 positions in at most 4 relevant lanes). A reasonable number of longitudinal goals of the first type is bounded by $8 \times 2 \times 3 = 48$ (8 relevant cars, whether to be in front or behind them, and 3 relevant distances). A reasonable number of absolute speed targets are 10, and a reasonable upper bound on the number of speed constraints is 2. To implement a given lateral or longitudinal goal, we need to apply acceleration and then deceleration (or the other way around). The
 10 aggressiveness of achieving the goal is a maximal (in absolute value) acceleration/deceleration to achieve the goal. With the goal and aggressiveness defined, we have a closed form formula to implement the goal, using kinematic calculations. The only remaining part is to determine the combination between the lateral and longitudinal goals (e.g., “start with the lateral goal, and exactly at the middle of it, start to apply also the longitudinal goal”). A set of 5 mixing times and 3 aggressiveness levels seems more than
 15 enough. All in all, we have obtained a semantic action space whose size is $\approx 10^4$.

[0503] It is worth mentioning that the variable time required for fulfilling these semantic actions is not the same as the frequency of the decision-making process. To be reactive to the dynamic world, we should make decisions at a high frequency — in our implementation, every 100 ms. In contrast, each such decision is based on constructing a trajectory that fulfills some semantic action, which
 20 will have a much longer time horizon (say, 10 seconds). We use the longer time horizon since it helps us to better evaluate the short-term prefix of the trajectory. The next subsection discusses the evaluation of semantic actions, but before that, we argue that semantic actions induce a sufficient search space.

[0504] As discussed above, a semantic action space induces a subset of all possible geometrical curves, whose size is exponentially smaller (in T) than enumerating all possible geometrical
 25 curves. The first immediate question is whether the set of short-term prefixes of this smaller search space contains all geometric commands that we will ever want to use. This is indeed sufficient in the following sense. If the road is free of other agents, then there is no reason to make changes except setting a lateral goal and/or absolute acceleration commands and/or speed constraints on certain positions. If the road contains other agents, we may want to negotiate the right of way with the other agents. In this case, it
 30 suffices to set longitudinal goals relatively to the other agents. The exact implementation of these goals in the long run may vary, but the short-term prefixes will not change by much. Hence, we obtain a very good cover of the relevant short-term geometrical commands.

[0505] Constructing an evaluation function for semantic actions

[0506] We have defined a semantic set of actions, denoted by A^s . Given that we are currently
 35 in state, s , we need a way to choose the best $a^s \in A^s$. To tackle this problem, we follow a similar approach to the *options mechanism* of [6]. The basic idea is to think of a^s as a meta-action (or an option). For each choice of a meta-action, we construct a geometrical trajectory $(s_1, a_1), \dots, (s_T, a_T)$ that represents an implementation of the meta-action, a^s . To do so we of course need to know how other agents will react to

our actions, but for now we are still relying on (the non-realistic) assumption that $s_{t+1} = f(s_t, a_t)$ for some known deterministic function f . We can now use $\frac{1}{T} \sum_{t=1}^T \rho(s_t, a_t)$ as a good approximation of the quality of performing the semantic action a^s when we are at state s^1 .

[0507] This approach can yield a powerful driving policy. However, in some situations a more
5 sophisticated quality function may be needed. For example, suppose that we are following a slow truck
before an exit lane, where we need to take the exit lane. One semantic option is to keep driving slowly
behind the truck. Another one is to overtake the truck, hoping that later we can get back to the exit lane
and make the exit on time. The quality measure described previously does not consider what will happen
after we will overtake the truck, and hence we will not choose the second semantic action even if there is
10 enough time to make the overtake and return to the exit lane. Machine learning can help us to construct a
better evaluation of semantic actions that will take into account more than the immediate semantic
actions. As previously discussed, learning a Q function over immediate geometric actions is problematic
due to the low signal-to-noise ratio (the lack of advantage). This is not problematic when considering
semantic actions, both because there is a large difference between performing the different semantic
15 actions and because the semantic time horizon (how many semantic actions we take into account) is very
small (probably less than three in most cases).

[0508] Another potential advantage of applying machine learning is for the sake of
generalization: we may set an adequate evaluation function for every road, by a manual inspection of the
properties of the road, which may involve some trial and error. Here, a machine learning approach, as
20 discussed above, can be trained on a large variety of road types so as to generalize to unseen roads as
well. The semantic action space according to the disclosed embodiments may allow for potential
benefits: semantic actions contain information on a long-time horizon, hence we can obtain a very
accurate evaluation of their quality while being resource efficient.

[0509] The dynamics of the other agents

[0510] So far, we have relied on the assumption that s_{t+1} is a deterministic function of s_t and a_t .
As emphasized previously, this assumption is not completely realistic as our actions affect the behavior of
other road users. While we do take into account some reactions of other agents to our actions (for
example, we assume that if we will perform a safe cut-in than the car behind us will adjust its speed so as
not to hit us from behind), it is not realistic to assume that we model all of the dynamics of other agents.
30 The solution to this problem is to re-apply the decision making at a high frequency, and by doing this,
constantly adapt our policy to the parts of the environment that are beyond our modeling. In a sense, one
can think of this as a Markovization of the world at every step.

[0511] Sensing

[0512] This section describes the sensing state, which is a description of the relevant
35 information of the scene, and forms the input to the driving policy module. By and large, the sensing
state contains static and dynamic objects. The static objects are lanes, physical road delimiters,
constraints on speed, constraints on the right of way, and information on occluders (e.g., a fence that
occludes relevant part of a merging road). Dynamic objects are vehicles (e.g., bounding box, speed,

acceleration), pedestrians (bounding box, speed, acceleration), traffic lights, dynamic road delimiters (e.g., cones at a construction area), temporary traffic signs and police activity, and other obstacles on the road (e.g., an animal, a mattress that fell from a truck, etc.).

[0513] In any reasonable sensor setting, we cannot expect to obtain the exact sensing state, s .

5 Instead, we view raw sensor and mapping data, which we denote by $x \in X$, and there is a sensing system that takes x and produces an approximate sensing state.

[0514] **Definition 16 (Sensing system)** Let S denotes the domain of sensing state and let X be the domain of raw sensor and mapping data. A sensing system is a function $\hat{s} : X \rightarrow S$.

[0515] It is important to understand when we should accept $\hat{s}(x)$ as a reasonable approximation
10 to s . The ultimate way to answer this question is by examining the implications of this approximation on the performance of our driving policy in general, and on the safety in particular. Following our safety-comfort distinction, here again we distinguish between sensing mistakes that lead to non-safe behavior and sensing mistakes that affect the comfort aspects of the ride. Before addressing the details, the type of errors a sensing system might make include:

- 15
- False negative: the sensing system misses an object
 - False positive: the sensing system indicates a “ghost” object
 - Inaccurate measurements: the sensing system correctly detects an object but incorrectly estimates its position or speed
 - Inaccurate semantic: the sensing system correctly detects an object but misinterpret its
20 semantic meaning, for example, the color of a traffic light

[0516] Comfort

[0517] Recall that for a semantic action a , we have used $Q(s, a)$ to denote our evaluation of a given that the current sensing state is s . Our policy picks the action $\pi(s) = \operatorname{argmax}_a Q(s, a)$. If we inject $\hat{s}(x)$ instead of s then the selected semantic action would be $\pi(\hat{s}(x)) = \operatorname{argmax}_a Q(\hat{s}(x), a)$. If $\pi(\hat{s}(x)) = \pi(s)$
25 then $\hat{s}(x)$ should be accepted as a good approximation to s . But, it is also not bad at all to pick $\pi(\hat{s}(x))$ as long as the quality of $\pi(\hat{s}(x))$ w.r.t. the true state, s , is almost optimal, namely, $Q(s, \pi(\hat{s}(x))) \geq Q(s, \pi(s)) - \epsilon$, for some parameter ϵ . We say that \hat{s} is ϵ -accurate w.r.t. Q in such case. Naturally, we cannot expect the sensing system to be ϵ -accurate all the time. We therefore also allow the sensing system to fail with some small probability δ . In such a case we say that \hat{s} is Probably (w.p. of at least $1 - \delta$), Approximately
30 (up to ϵ), Correct, or PAC for short (borrowing Valiant’s PAC learning terminology).

[0518] We may use several (ϵ, δ) pairs for evaluating different aspects of the system. For example, we can choose three thresholds, $\epsilon_1 < \epsilon_2 < \epsilon_3$ to represent mild, medium, and gross mistakes, and for each one of them set a different value of δ . This leads to the following definition.

[0519] **Definition 17 (PAC sensing system)** Let $((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k))$ be a set of
35 (accuracy, confidence) pairs, let S be the sensing state domain, let X be the raw sensor and mapping data domain, and let D be a distribution over $X \times S$. Let A be an action space, $Q : S \times A \rightarrow \mathbb{R}$ be a quality function, and $\pi : S \rightarrow A$ be such that $\pi(s) \in \operatorname{argmax}_a Q(s, a)$. A sensing system, $\hat{s} : X \rightarrow S$, is Probably-

Approximately-Correct (PAC) with respect to the above parameters if for every $i \in \{1, \dots, k\}$ we have that

$$P_{(x,s) \sim D}[Q(s, \pi(\hat{s}(x))) \geq Q(s, \pi(s)) - \epsilon_i] \geq 1 - \delta_i.$$

[0520] Here, the definition depends on a distribution D over $X \times S$. It is important to emphasize that we construct this distribution by recording data of many human drivers but not by following the particular policy of our autonomous vehicle. While the latter seems more adequate, it necessitates online validation, which makes the development of the sensing system impractical. Since the effect of any reasonable policy on D is minor, by applying simple data augmentation techniques we can construct an adequate distribution and then perform offline validation after every major update of the sensing system. The definition provides a sufficient, but not necessary, condition for comfort ride using \hat{s} . It is not necessary because it ignores the important fact that short term wrong decisions have little effect on the comfort of the ride. For example, suppose that there is a vehicle 100 meters ahead, and it is slower than the host vehicle. The best decision would be to start accelerating slightly now. If the sensing system misses this vehicle but will detect it in the next time (after 100 milli-seconds), then the difference between the two rides will not be noticeable. To simplify the presentation, we have neglected this issue and required a stronger condition. The adaptation to a multi-frame PAC definition is conceptually straightforward but is more technical.

[0521] We next derive design principles that follow from the above PAC definition. Recall that we have described several types of sensing mistakes. For mistakes of types false negative, false positive, and inaccurate semantic, either the mistakes will be on non-relevant objects (e.g., a traffic light for left turn when we are proceeding straight), or they will be captured by the δ part of the definition. We therefore focus on the “inaccurate measurements” type of errors, which happens frequently.

[0522] Somewhat surprisingly, we will show that the popular approach of measuring the accuracy of a sensing system via ego-accuracy (that is, by measuring the accuracy of position of every object with respect to the host vehicle) is not sufficient for ensuring PAC sensing system. We will then propose a different approach that ensures PAC sensing system and will show how to obtain it efficiently. We start with some additional definitions.

[0523] For every object o in the scene, let $p(o)$, $\hat{p}(o)$ be the positions of o in the coordinate system of the host vehicle according to s , $\hat{s}(x)$, respectively. Note that the distance between o and the host vehicle is $\|p\|$. The *additive* error of \hat{p} is $\|p(o) - \hat{p}(o)\|$. The *relative* error of $\hat{p}(o)$, w.r.t. the distance between o and the host vehicle, is the additive error divided by $\|p(o)\|$, namely $\frac{\|p(o) - \hat{p}(o)\|}{\|p(o)\|}$.

[0524] It is not realistic to require that the additive error is small for far away objects. Indeed, consider o to be a vehicle at a distance of 150 meters from the host vehicle, and let ϵ be of moderate size, say $\epsilon = 0.1$. For additive accuracy, it means that we should know the position of the vehicle up to 10 cm of accuracy. This is not realistic for reasonably priced sensors. On the other hand, for relative accuracy we need to estimate the position up to 10%, which amounts to 15 m of accuracy. This is feasible to achieve (as described below).

[0525] A sensing system, \hat{s} , positions a set of objects, O , in an ϵ -ego-accurate way, if for every $o \in O$, the (relative) error between $p(o)$ and $\hat{p}(o)$ is at most ϵ . The following example demonstrates that

an ϵ -ego-accurate sensing state does not guarantee PAC sensing system with respect to every reasonable Q . Indeed, consider a scenario in which the host vehicle drives at a speed of 30 m/s, and there is a stopped vehicle 150 meters in front of it. If this vehicle is in the ego lane, and there is no option to change lanes in time, we must start decelerating now at a rate of at least $3m/s^2$ (otherwise, we will either
 5 not stop in time or we will need to decelerate strongly later). On the other hand, if the vehicle is on the side of the road, we don't need to apply a strong deceleration. Suppose that $p(o)$ is one of these cases while $\hat{p}(o)$ is the other case, and there is a 5 meters difference between these two positions. Then, the relative error of $\hat{p}(o)$ is

$$\frac{\|\hat{p}(o) - p(o)\|}{\|p(o)\|} = \frac{5}{150} = \frac{1}{30} \leq 0.034.$$

10 [0526] That is, the sensing system may be ϵ -ego-accurate for a rather small value of ϵ (less than 3.5% error), and yet, for any reasonable Q function, the values of Q are completely different since we are confusing between a situation in which we need to brake strongly and a situation in which we do not need to brake strongly.

[0527] The above example shows that ϵ -ego-accuracy does not guarantee that our sensing
 15 system is PAC. Whether there is another property that is sufficient for PAC sensing system depends on Q . We will describe a family of Q functions for which there is a simple property of the positioning that guarantees PAC sensing system. The problem of ϵ -ego-accuracy is that it might lead to semantic mistakes — in the aforementioned example, even though \hat{s} was ϵ -ego-accurate with $\epsilon < 3.5\%$, it mis-assigned the vehicle to the correct lane. To solve this problem, we rely on semantic units for lateral
 20 position.

[0528] **Definition 18 (semantic units)** *A lane center is a simple natural curve, namely, it is a differentiable, injective, mapping $\ell : [a, b] \rightarrow \mathbb{R}^3$, where for every $a \leq t_1 < t_2 \leq b$ we have that the length $\text{Length}(t_1, t_2) := \int_{\tau=t_1}^{t_2} |\ell'(\tau)| d\tau$ equals to $t_2 - t_1$. The width of the lane is a function $w : [a, b] \rightarrow \mathbb{R}_+$. The projection of a point $x \in \mathbb{R}^3$ onto the curve is the point on the curve closest to x , namely, the point $\ell(t_x)$ for
 25 $t_x = \arg\min_{t \in [a, b]} \|\ell(t) - x\|$. The semantic longitudinal position of x w.r.t. the lane is t_x and the semantic lateral position of x w.r.t. the lane is $\ell(t_x)/w(t_x)$. Semantic speed and acceleration are defined as first and second derivatives of the above.*

[0529] Similarly to geometrical units, for semantic longitudinal distance we use relative error: if \hat{s} induces a semantic longitudinal distance of $\hat{p}(o)$ for some object, while the true distance is $p(o)$, then
 30 the relative error is $\frac{|\hat{p}(o) - p(o)|}{\max\{p(o), 1\}}$ (where the maximum in the denominator deals with cases in which the object has almost the same longitudinal distance (e.g., a car next to us on another lane). Since semantic lateral distances are small we can use additive error for them. This leads to the following definition:

[0530] **Definition 19 (error in semantic units)** *Let ℓ be a lane and suppose that the semantic longitudinal distance of the host vehicle w.r.t. the lane is 0. Let $x \in \mathbb{R}^3$ be a point and let $p_{\text{lat}}(x)$, $p_{\text{lon}}(x)$
 35 be the semantic lateral and longitudinal distances to the point w.r.t. the lane. Let $\hat{p}_{\text{lat}}(x)$, $\hat{p}_{\text{lon}}(x)$ be approximated measurements. The distance between \hat{p} and p w.r.t. x is defined as*

$$d(\hat{p}, p; x) = \max \left\{ |\hat{p}_{\text{lat}}(x) - p_{\text{lat}}(x)|, \frac{|\hat{p}_{\text{lon}}(x) - p_{\text{lon}}(x)|}{\max\{p_{\text{lon}}(x), 1\}} \right\}$$

[0531] The distance of the lateral and longitudinal velocities is defined analogously. Equipped with the above definition, we are ready to define the property of Q and the corresponding sufficient condition for PAC sensing system.

5 [0532] **Definition 20 (Semantically-Lipschitz Q)** A Q function is L -semantically-Lipschitz if for every a, s, \hat{s} , $|Q(s, a) - Q(\hat{s}(x), a)| \leq L \max_o d(\hat{p}, p; o)$, where \hat{p}, p are the measurements induced by s, \hat{s} on an object o .

[0533] As an immediate corollary we obtain:

[0534] **Lemma 8** If Q is L -semantically-Lipschitz and a sensing system \hat{s} produces semantic
10 measurements such that with probability of at least $1 - \delta$ we have $d(\hat{p}, p; o) \leq \theta/L$, then \hat{s} is a PAC sensing system with parameters θ, δ .

[0535] Safety

[0536] This section discusses the potential for sensing errors that can lead to unwanted behaviors. As mentioned before, the policy is provably safe, in the sense that it won't lead to accidents of
15 the host AV's blame. Such accidents might still occur due to hardware failure (e.g., a breakdown of all the sensors or exploding tire on the highway), software failure (a significant bug in some of the modules), or a sensing mistake. Our ultimate goal is that the probability of such events will be extremely small — a probability of 10^{-9} for such an accident per hour. To appreciate this number, the average number of hours driver in the U.S. spends on the road is (as of 2016) less than 300 hours. So, in expectation, one would
20 need to live 3.3 million years to be in an accident resulting from one of these types of events.

[0537] We first define what is a safety-related sensing error. Recall that at every step, our policy picks the value of a that maximizes $Q(s, a)$, namely, $\pi(s) = \operatorname{argmax}_a Q(s, a)$. We ensure safety by letting $Q(s, a) = -\infty$ for every action a that is not cautious (see Definition 14). Therefore, the first type of safety-critic sensing mistake is if our sensing system leads to picking a non-safe action. Formally, letting
25 $\pi(\hat{s}(x)) = \operatorname{argmax}_a Q(\hat{s}(x), a)$ be the decision according to \hat{s} , we say that \hat{s} leads to a *safety-critic miss* if $Q(s, \pi(\hat{s}(x))) = -\infty$. The second type of safety-critic sensing mistake is if all the actions are non-safe according to $\hat{s}(x)$, and we must apply the standard emergency policy (e.g., braking hard), while according to s there is a safe action, namely, $\max_a Q(s, a) > -\infty$. This is dangerous when our speed is high and there is a car behind us. We call such mistake a *safety-critic ghost*.

[0538] Usually, a safety-critic miss is caused by a false negative while a safety-critic ghost is caused by a false positive. Such mistakes can also be caused from significantly incorrect measurements, but in most cases, our comfort objective ensures we are far away from the boundary of the safety definitions, and therefore reasonable measurements errors are unlikely to lead to a safety-critic mistake. How can we ensure that the probability of safety-critic mistakes will be very small, say, smaller than 10^{-9}
35 per hour? As followed from Lemma 1, without making further assumptions we need to check our system on more than 10^9 hours of driving. This is unrealistic (or at least extremely challenging) — it amounts to recording the driving of 3.3 million cars over a year. Furthermore, building a system that achieves such a

high accuracy is a great challenge. A solution for both the system design and validation challenges is to rely on several sub-systems, each of which is engineered independently and depends on a different technology, and the systems are fused together in a way that ensures boosting of their individual accuracy.

5 [0539] Suppose we build three sub-systems, denoted, s_1, s_2, s_3 (the extension to more than 3 is straightforward). Each sub-system receives a and should output safe/non-safe. Actions for which the majority of the sub-systems (2 in our case) accept as safe are considered safe. If there is no action that is considered safe by at least 2 sub-systems, then the default emergency policy is applied. The performance of this fusion scheme is analyzed as follows based on the following definition:

10 [0540] **Definition 21 (One side c -approximate independent)** *Two Bernoulli random variables r_1, r_2 are called one side c -approximate independent if*

$$\mathbb{P}[r_1 \wedge r_2] \leq c \mathbb{P}[r_1] \mathbb{P}[r_2].$$

[0541] For $i \in \{1,2,3\}$, denote by e_i^m, e_i^g the Bernoulli random variables that indicate if sub-system i performs a safety-critic miss/ghost respectively. Similarly, e^m, e^g indicate a safety-critic miss/ghost of the fusion system. We rely on the assumption that for any pair $i \neq j$, the random variables e_i^m, e_j^m are one sided c -approximate independent, and the same holds for e_i^g, e_j^g . Before explaining why this assumption is reasonable, let us first analyze its implication. We can bound the probability of e^m by:

$$\begin{aligned} \mathbb{P}[e^m] &= \mathbb{P}[e_1^m \wedge e_2^m \wedge e_3^m] + \sum_{j=1}^3 \mathbb{P}[\neg e_j^m \wedge \bigwedge_{i \neq j} e_i^m] \\ &\leq 3\mathbb{P}[e_1^m \wedge e_2^m \wedge e_3^m] + \sum_{j=1}^3 \mathbb{P}[\neg e_j^m \wedge \bigwedge_{i \neq j} e_i^m] \\ &= \sum_{j=1}^3 \mathbb{P}[\bigwedge_{i \neq j} e_i^m] \\ &\leq c \sum_{j=1}^3 \prod_{i \neq j} \mathbb{P}[e_i^m]. \end{aligned}$$

20

[0542] Therefore, if all sub-systems have $\mathbb{P}[e_i^m] \leq p$ then $\mathbb{P}[e^m] \leq 3cp^2$. The exact same derivation holds for the safety-critic ghost mistakes. By applying a union bound we therefore conclude:

[0543] **Corollary 2** *Assume that for any pair $i \neq j$, the random variables e_i^m, e_j^m are one sided c -approximate independent, and the same holds for e_i^g, e_j^g . Assume also that for every i , $\mathbb{P}[e_i^m] \leq p$ and $\mathbb{P}[e_i^g] \leq p$. Then,*

25

$$\mathbb{P}[e^m \vee e^g] \leq 6cp^2.$$

[0544] This corollary allows us to use significantly smaller data sets in order to validate the sensing system. For example, if we would like to achieve a safety-critic mistake probability of 10^{-9} , instead of taking order of 10^9 examples, it suffices to take order of 10^5 examples and test each system separately.

30

[0545] There may be pairs of sensors that yield non-correlated errors. For example, radar works well in bad weather conditions but might fail due to non-relevant metallic objects, as opposed to camera that is affected by bad weather but is not likely to be affected by metallic objects. Seemingly, camera and lidar have common sources of error — e.g., both are affected by foggy weather, heavy rain, or snow. However, the type of errors for camera and lidar would be different — a camera might miss objects due to bad weather and lidar might detect a ghost due to reflections from particles in the air. Since we have distinguished between the two types of errors, the approximate independency is still likely to hold.

[0546] Our definition of safety-important ghost requires that *all* actions are non-safe by at least two sensors. Even in difficult conditions (e.g., heavy fog), this is unlikely to happen. The reason is that in such situations, systems that are affected by the difficult conditions (e.g., the lidar), will dictate a very defensive driving, as they can declare high velocity and lateral maneuver to be non-safe actions. As a result, the host AV will drive slowly, and then even if an emergency stop is required, it is not dangerous due to the low speed of driving. Therefore, our definition yields an adaptation of the driving style to the conditions of the road.

[0547] Building a scalable sensing system

[0548] The requirements from a sensing system, both in terms of comfort and safety, have been described. Next, an approach for building a sensing system that meets these requirements while being scalable is described. There are three main components of the sensing system. The first is long range, 360 degrees coverage, of the scene based on cameras. The three main advantages of cameras are: (1) high resolution, (2) texture, (3) price. The low price enables a scalable system. The texture enables to understand the semantics of the scene, including lane marks, traffic light, intentions of pedestrians, and more. The high resolution enables a long range of detection. Furthermore, detecting lane marks and objects in the same domain enables excellent semantic lateral accuracy. The two main disadvantages of cameras are: (1) the information is 2D and estimating longitudinal distance is difficult, (2) sensitivity to lighting conditions (low sun, bad weather). We overcome these difficulties using the next two components of our system.

[0549] The second component of our system is a semantic high-definition mapping technology, called Road Experience Management (REM) (which involves navigation based on target trajectories predetermined and stored for road segments along with an ability to determine precise locations along the target trajectories based on the location (e.g., in images) of recognized landmarks identified in the environment of the host vehicle). A common geometrical approach to map creation is to record a cloud of 3D points (obtained by a lidar) in the map creation process, and then, localization on the map is obtained by matching the existing lidar points to the ones in the map. There are several disadvantages of this approach. First, it requires a large memory per kilometer of mapping data, as we need to save many points. This necessitates an expensive communication infrastructure. Second, not all cars may be equipped with lidar sensors, and therefore, the map is updated very infrequently. This is problematic as changes in the road can occur (construction zones, hazards), and the “time-to-reflect-

reality” of lidar-based mapping solutions is large. In contrast, REM follows a semantic-based approach. The idea is to leverage the large number of vehicles that are equipped with cameras and with software that detects semantically meaningful objects in the scene (lane marks, curbs, poles, traffic lights, etc.). Nowadays, many new cars are equipped with ADAS systems which can be leveraged for crowd-sourced map creation. Since the processing is done on the vehicle side, only a small amount of semantic data should be communicated to the cloud. This allows a very frequent update of the map in a scalable way. In addition, the autonomous vehicles can receive the small sized mapping data over existing communication platforms (the cellular network). Finally, highly accurate localization on the map can be obtained based on cameras, without the need for expensive lidars.

[0550] REM may be used for several purposes. First, it gives us a foresight on the static structure of the road (we can plan for a highway exit way in advance). Second, it gives us another source of accurate information of all of the static information, which together with the camera detections yields a robust view of the static part of the world. Third, it solves the problem of lifting the 2D information from the image plane into the 3D world as follows. The map describes all of the lanes as curves in the 3D world. Localization of the ego vehicle on the map enables to trivially lift every object on the road from the image plane to its 3D position. This yields a positioning system that adheres to the accuracy in semantic units. A third component of the system may be complementary radar and lidar systems. These systems may serve two purposes. First, they can offer extremely high levels of accuracy for augmenting safety. Second, they can give direct measurements on speed and distances, which further improves the comfort of the ride.

[0551] The following sections include technical lemmas and several practical considerations of the RSS system.

[0552] **Lemma 9** For all $x \in [0, 0.1]$, it holds that $1 - x > e^{-2x}$.

Proof Let $f(x) = 1 - x - e^{-2x}$. Our goal is to show it is ≥ 0 for $x \in [0, 0.1]$. Note that $f(0) = 0$, and it is therefore sufficient to have that $f(x) \geq 0$ in the aforementioned range. Explicitly, $f'(x) = -1 + 2e^{-2x}$. Clearly, $f'(0) = 1$, and it is monotonically decreasing, hence it is sufficient to verify that $f'(0.1) > 0$, which is easy to do numerically, $f'(0.1) \approx 0.637$.

[0553] Efficient Cautiousness Verification - Occluded Objects

[0554] As with non-occluded objects, we can check whether giving the current command, and after it, the DEP, is RSS. For this, we unroll our future until t_{brake} , when assuming the exposure time is 1 and we then command DEP, which suffices for cautiousness by definition. For all $t' \in [0, t_{\text{brake}}]$, we check whether a blameful accident can occur - when assuming worst case over the occluded object. We use some of our worst case maneuvers and safe distance rules. We take an occluder based approach to find interest points - namely, for each occluding object, we calculate the worst case. This is a crucial efficiency-driven approach - a pedestrian, for example, can be hidden in many positions behind a car, and can perform many maneuvers, but there’s a single worst case position and maneuver it can perform.

[0555] Next, consider the more elaborate case, that of the occluded pedestrian. Consider an occluded area behind a parked car. The closest points in an occluded area and the front/side of our car c

may be found, for example, by their geometrical properties (triangles, rectangles). Formally, we can consider the occluded area as a union of a small number of convex regions of simple shape, and treat each of them separately. Furthermore, it can be seen that a pedestrian can run into the front of the car (under the v_{limit} constraint) from the occluded area IFF he can do it using the shortest path possible. Using the fact that the maximal distance which can be travelled by the pedestrian is $v_{limit} \cdot t'$, we obtain a simple check for a frontal hit possibility. As for a side hit, we note that in case the path is shorter than $v_{limit} \cdot t'$, we are responsible IFF our lateral velocity is greater than μ , in the direction of the hit. Disclosed is an algorithm for cautiousness verification with respect to occluded pedestrians, which is described here in free pseudo code. The crucial part, that of checking existence of possibility of blameful accident with a pedestrian occluded by a vehicle, is done in the simple manner described above.

```

Algorithm 2: Check cautiousness w.r.t. an occluded pedestrian

for  $t' \in [0, t_{brake}]$ 
  Roll self future until  $t'$ 
  if Exists possibility of blameful accident with a pedestrian occluded by a vehicle
    return "non-cautious"
  return "cautious"

```

[0556] On the Problem of Validating a Simulator

[0557] As previously discussed, multi-agent safety may be difficult to validate statistically as it should be done in an "online" manner. One may argue that by building a simulator of the driving environment, we can validate the driving policy in the "lab." However, validating that the simulator faithfully represents reality is as hard as validating the policy itself. To see why this is true, suppose that the simulator has been validated in the sense that applying a driving policy π in the simulator leads to a probability of an accident of \hat{p} , and the probability of an accident of π in the real world is p , with $|p - \hat{p}| < 0$. (We need that 0 will be smaller than 10^{-9} .) Next replace the driving policy to be π' . Suppose that with probability of 10^{-8} , π' performs a weird action that confuses human drivers and leads to an accident. It is possible (and even rather likely) that this weird action is not modeled in the simulator, without contradicting its superb capabilities in estimating the performance of the original policy π . This proves that even if a simulator has been shown to reflect reality for a driving policy π , it is not guaranteed to reflect reality for another driving policy.

[0558] The Lane-Based Coordinate System

[0559] One simplifying assumption that can be made in the RSS definition is that the road is comprised by adjacent, straight lanes, of constant width. The distinction between lateral and longitudinal axes, along with an ordering of longitudinal position, may play a significant role in RSS. Moreover, the definition of those directions is clearly based on the lane shape. A transformation from (global) positions on the plane, to a lane-based coordinate system reduces the problem to the original, "straight lane of constant width," case.

[0560] Assume that the lane's center is a smooth directed curve r on the plane, where all of its pieces, denoted $r^{(1)}, \dots, r^{(k)}$, are either linear, or an arc. Note that smoothness of the curve implies that no pair of consecutive pieces can be linear. Formally, the curve maps a "longitudinal" parameter, $Y \in [Y_{\min}, Y_{\max}] \subset \mathbb{R}$, into the plane, namely, the curve is a function of the form $r : [Y_{\min}, Y_{\max}] \rightarrow \mathbb{R}^2$. We define a continuous lane-width function $w : [Y_{\min}, Y_{\max}] \rightarrow \mathbb{R}_+$, mapping the longitudinal position Y into a positive lane width value. For each Y , from smoothness of r , we can define the normal unit-vector to the curve at position Y , denoted $r^\perp(Y)$. We naturally define the subset of points on the plane which reside in the lane as follows:

$$R = \{r(Y) + \alpha w(Y)r^\perp(Y) \mid Y \in [Y_{\min}, Y_{\max}], \alpha \in [-1/2, 1/2]\}$$

[0561] Informally, our goal is to construct a transformation ϕ of R into such that for two cars which are on the lane, their "logical ordering" will be preserved: if c_r is "behind" c_f on the curve, then $\phi(c_r)_y < \phi(c_f)_y$. If c_l is "to the left of" c_r on the curve, then $\phi(c_l)_x < \phi(c_r)_x$. Where, as in RSS, we will associate the y -axis with the "longitudinal" axis, and the x -axis with the "lateral".

[0562] To define ϕ , we rely on the assumption that for all i , if $r^{(i)}$ is an arc of radius ρ , then the width of the lane throughout $r^{(i)}$ is $\leq \rho/2$. Note that this assumption holds for any practical road. The assumption trivially implies that for all $(x', y') \in R$, there exists a unique pair $Y' \in [Y_{\min}, Y_{\max}], \alpha' \in [-1/2, 1/2]$, s.t. $(x', y') = r(Y') + \alpha' w(Y') r^\perp(Y')$. We can now define $\phi : R \rightarrow \mathbb{R}^2$ to be $\phi(x', y') = (Y', \alpha')$, where (Y', α') are the unique values that satisfy $(x', y') = r(Y') + \alpha' w(Y') r^\perp(Y')$.

[0563] This definition captures the notion of a "lateral maneuver" in lane's coordinate system. Consider, for example, a widening lane, with a car driving exactly on one of the lane's boundaries (see Fig. 23). The widening of lane 2301 means that the car 2303 is moving away from the center of the lane, and therefore has lateral velocity with respect to the lane. However, this doesn't mean it performs a lateral maneuver. The definition of $\phi(x', y')_x = \alpha'$, namely, the lateral distance to the lane's center in $w(Y')$ -units, implies that the lane boundaries have a fixed lateral position of $\pm 1/2$. Hence, a car that sticks to one of the lane's boundaries is not considered to perform any lateral movement. Finally, it can be seen that ϕ is a homomorphism. The term *lane-based coordinate system* is used when discussing $\phi(R) = [Y_{\min}, Y_{\max}] \times [-1/2, 1/2]$. We have thus obtained a reduction from a general lane geometry to a straight, longitudinal/lateral, coordinate system.

[0564] Extending RSS to General Road Structure

[0565] In this section a complete definition of RSS that holds for all road structures is described. This section deals with the definition of RSS and not on how to efficiently ensure that a policy adheres to RSS. The concept of *route priority* is next introduced to capture any situation in which more than a single lane geometry exists, for example junctions.

[0566] The second generalization deals with two-way roads, in which there can be two cars driving at opposite directions. For this case, the already established RSS definition is still valid, with the minor generalization of "safe distance" to oncoming traffic. Controlled junctions (that use traffic lights to dictate the flow of traffic), may be fully handled by the concepts of route priority and two-way roads. Unstructured roads (for example parking areas), where there is no clear route definition may also be

handled with RSS. RSS is still valid for this case, where the only needed modification is a way to define virtual routes and to assign each car to (possibly several) routes.

[0567] Route Priority

[0568] The concept of *route priority* is now introduced to deal with scenarios in which there are multiple different road geometries in one scene that overlap in a certain area. Examples, as shown in Figs. 24A-D, include roundabouts, junctions, and merge into highways. A way to transform general lane geometry into a lane-based one, with coherent meaning for longitudinal and lateral axes, has been described. Now scenarios in which multiple routes of different road geometry exists are addressed. It follows that when two vehicles approach the overlap area, both perform a cut-in to the frontal corridor of the other one. This phenomenon cannot happen when two routes have the same geometry (as is the case of two adjacent highway lanes). Roughly speaking, the principle of route priority states that if routes r_1 , r_2 overlap, and r_1 has priority over r_2 , then a vehicle coming from r_1 that enters into the frontal corridor of a vehicle that comes from r_2 is not considered to perform a cut-in.

[0569] To explain the concept formally, recall that the blame of an accident depends on geometrical properties which are derived from the lane's coordinate system, and on worst case assumptions which rely on it too. Let r_1, \dots, r_k be the routes defining the road's structure. As a simple example, consider the merge scenario, as depicted in Fig. 24A. Assume two cars, 2401 (c_1) and 2402 (c_2) are driving on routes r_1, r_2 respectively, and r_1 is the prioritized route. For example, suppose that r_1 is a highway lane and r_2 is a merging lane. Having defined the route-based coordinate systems for each route, a first observation is that we can consider any maneuver in any route's coordinate system. For example, if we use r_2 's coordinate system, driving straight on r_1 seems like a merge into r_2 's left side. One approach to definition of RSS could be that each of the cars can perform a maneuver IFF $\forall i \in \{1, 2\}$, if it is safe with respect to r_i . However, this implies that c_1 , driving on the prioritized route, should be very conservative w.r.t. r_2 , the merging route, as c_2 can drive exactly on the route, and hence can win by lateral position. This is unnatural, as cars on the highway have the right-of-way in this case. To overcome this problem, we define certain areas in which route priority is defined, and only some of the routes are considered as relevant for safety.

[0570] **Definition 22 (Accident Responsibility with Route Priority)** Suppose r_1, r_2 are two routes with different geometry that overlap. We use $r_1 \succ_{[b,e]} r_2$ to symbolize that r_1 has priority over r_2 in the longitudinal interval $[b, e]$ (Fig. 25) of r_1 coordinate system. Suppose there is an accident between cars c_1, c_2 , driving on routes r_1, r_2 . For $i \in \{1, 2\}$, let $b_i \subset \{1, 2\}$ indicate the cars to blame for the accident if we consider the coordinate system of r_i . The blame for the accident is as follows:

- If $r_1 \succ_{[b,e]} r_2$ and on the blame time w.r.t. r_1 , one of the cars was in the interval $[b, e]$ of the r_1 -system's longitudinal axis, then the blame is according to b_1 .
- Otherwise, the blame is according to $b_1 \cup b_2$.

[0571] To illustrate the definition, consider again the merge into highway example. The lines denoted "b" and "e" in Fig. 25 indicate the values of b, e for which $r_1 \succ_{[b,e]} r_2$. Thus, we allow cars to drive naturally on the highway, while implying merging cars must be safe with respect to those cars. In

particular, observe that in the case a car 2401 c_1 drives at the center of the prioritized lane, with no lateral velocity, it will not be blamed for an accident with a car 2402 c_2 driving on a non-prioritized lane, unless 2402 c_2 has cut-in into 2401 c_1 's corridor at a safe distance. Note, that the end result is very similar to the regular RSS - this is exactly the same as a case where a car, on a straight road, tries to perform a lane

5 change. Note that there may be cases where the route used by another agent is unknown. For example, in Fig. 26, car 2601 may not be able to determine whether a car 2602 will take path “a” or path “b”. In such cases, RSS may be obtained by iteratively checking all possibilities.

[0572] Two-Way Traffic

[0573] To deal with two-way traffic, the modification to the blame definition comes through

10 sharpening the parts which rely on rear/front relationships, as those are of a slightly different meaning in such cases. Consider two cars c_1, c_2 driving on some straight two lane road, in opposite longitudinal directions, namely, $v_{1,\text{long}} \cdot v_{2,\text{long}} < 0$. The direction of driving with respect to a lane may be negative in reasonable urban scenarios, such as a car deviating to the opposite lane to overtake a parked truck, or a car reversing into a parking spot. It is therefore required to extend the definition of the safe longitudinal

15 distance which we have introduced for cases where negative longitudinal velocity was assumed to be unrealistic. Recall that a distance between c_r, c_f was safe if a maximal brake by c_f would allow enough of response time for c_r to brake before crashing into c_f . In our case, we again consider the “worst-case” by the opposite car, in a slightly different manner: of course we do not assume that the “worst case” is that it speeds up towards us, but that it indeed will brake to avoid a crash - but only using some reasonable

20 braking power. In order to capture the difference in responsibility between the cars, when one of them clearly drives at the opposite direction, we start by defining a “correct” driving direction.

[0574] In the RSS definition for parallel lanes, the relevant lane has been defined as the one whose center is closest to the cut-in position. We can now reduce ourselves to consideration of this lane (or, in the case of symmetry, deal with the two lanes separately, as in Definition 22). In the definition

25 below, the term “heading” denotes the arc tangent (in radians) of the lateral velocity divided by the longitudinal velocity.

[0575] **Definition 23 ((μ_1, μ_2, μ_3)-Winning by Correct Driving Direction)** *Assume c_1, c_2 are driving in opposite directions, namely $v_{1,\text{long}} \cdot v_{2,\text{long}} < 0$. Let x_i, h_i be their lateral positions and headings, w.r.t. the lane. We say that c_1 (μ_1, μ_2, μ_3)-Wins by Correct Driving Direction if all of the following*

30 *conditions hold:*

- $|h_1| < \mu_1,$
- $|h_2 - \pi| < \mu_2,$
- $|x_1| < \mu_3.$

[0576] The indicator of this event is denoted by $W_{\text{CDD}}(i)$.

[0577] In words, c_1 wins if it drives close to the lane center, in the correct direction, while c_2 takes the opposite direction. At most one car can win, and it can be that none of the cars does so.

Intuitively, assume there is a crash in the discussed situation. It is reasonable to put more responsibility

over a car c_1 , that loses by Correct Driving Direction. This is done by re-defining $a_{\max, \text{brake}}$ for the case that a car wins by correct driving direction.

[0578] **Definition 24 (Reasonable Braking Power)** Let $\alpha_{\max, \text{brake}, \text{wedd}} > 0$ be a constant, smaller than $a_{\max, \text{brake}}$. Assume c_1, c_2 are driving in opposite directions. The Reasonable Braking Power of each car c_i , denoted RBP_i is $\alpha_{\max, \text{brake}, \text{wedd}}$ if c_i (μ_1, μ_2, μ_3)-Wins by Correct Driving Direction and
5 $a_{\max, \text{brake}}$ otherwise.

[0579] The exact values of $\alpha_{\max, \text{brake}, \text{wedd}}, a_{\max, \text{brake}}$, for the cases of winning/not-winning by correct driving direction, are constants which should be defined, and can depend on the type of road and the lanes driven by each car. For example, in a narrow urban street, it may be the case that winning by a
10 correct driving direction does not imply a much lesser brake value: in dense traffic, we do expect a car to brake at similar forces, either when someone clearly deviated into its lane or not. However, consider an example of a rural two way road, where high speeds are allowed. When deviating to the opposite lane, cars which drive at the correct direction cannot be expected to apply a very strong braking power to avoid hitting the host vehicle - the host vehicle will have more responsibility than them. Different constants can
15 be defined for the case when two cars are at the same lane, with one of them reversing into a parking spot.

[0580] The safety distance between cars which are driving in opposite directions, and immediately derive its exact value, is next defined.

[0581] **Definition 25 (Safe Longitudinal Distance - Two-Way Traffic)** A longitudinal distance between a car c_1 and another car c_2 which are driving in opposite directions and are both in the
20 frontal corridors of each other, is safe w.r.t. a response time ρ if for any acceleration command a , $|a| < a_{\max, \text{accel}}$, performed by c_1, c_2 until time ρ , if c_1 and c_2 will apply their Reasonable Braking Power from time ρ until a full stop then they won't collide.

[0582] **Lemma 10** Let c_1, c_2 as in Definition 25. Let $RBP_i, a_{\max, \text{accel}}$ be the reasonable braking (for each i) and acceleration commands, and let ρ be the cars' response time. Let v_1, v_2 be the
25 longitudinal velocities of the cars, and let l_1, l_2 be their lengths. Define $v_{i, \rho, \max} = |v_i| + \rho \cdot a_{\max, \text{accel}}$. Let $L = (l_r + l_f)/2$. Then, the minimal safe longitudinal distance is:

$$d_{\min} = L + \sum_{i=1}^2 \left(\frac{|v_i| + v_{i, \rho, \max}}{2} \rho + \frac{v_{i, \rho, \max}^2}{2RBP_i} \right)$$

[0583] It can be seen that the term in the sum is the maximal distance travelled by each car until it reaches full stop, when performing the maneuver from Definition 25. Therefore, in order for the
30 full stop to be at a distance greater than L , the initial distance must be larger than this sum and an additional term of L .

[0584] The same blame time definition of RSS, with the non-safe longitudinal distance as defined in Definition 25, is used to define the blame/accident responsibility for a two-way traffic scenario.

[0585] **Definition 26 (Blame in Two-Way Traffic)** The Blame in Two-Way Traffic of an
35 accident between cars c_1, c_2 driving in opposite directions, is a function of the state at the Blame Time, and is defined as follows:

- *If the Blame Time is also a cut-in time, the blame is defined as in the regular RSS definition.*
- *Otherwise, for every i , the blame is on c_i if at some t that happens after the blame time, c_i was not braking at a power of at least RBP_i .*

5 [0586] For example, assume a safe cut-in occurred before the blame time. For example, c_1 has deviated to the opposite lane, performed a cut-in into c_2 's corridor, at a safe distance. Note that c_2 wins by correct driving direction, and hence this distance can be very large - we do not expect c_2 to perform strong braking power, but only the Reasonable Braking Power. Then, both cars have responsibility not to crash into each other. However, if the cut-in was not in a safe-distance, we use the regular definition,
 10 noting that c_2 will not be blamed if it drove in the center of its lane, without lateral movement. The blame will be solely on c_1 . This allows a car to drive naturally at the center of its lane, without worrying about traffic which may unsafely deviate into its corridor. On the other hand, safe deviation to the opposite lane, a common maneuver required in dense urban traffic, is allowed. Considering the example of a car which initiates a reverse parking maneuver, it should start reversing while making sure the distance to
 15 cars behind it is safe.

[0587] Traffic Lights

[0588] In scenarios that include intersections with traffic lights, one might think that the simple rule for traffic lights scenarios is "if one car's route has the green light and the other car's route has a red light, then the blame is on the one whose route has the red light". However, this is not the
 20 correct rule, especially in all cases. Consider for example the scenario depicted in Fig. 27. Even if the car 2701 is on a route having a green light, we do not expect it to ignore car 2703 that is already in the intersection. The correct rule is that the route that has a green light has priority over routes that have a red light. Therefore, we obtain clear reduction from traffic lights to the route priority concept we have described previously.

25 [0589] Unstructured Road

[0590] Turning to roads where no clear route geometry can be defined, consider first a scenario where there is no lane structure at all (e.g. a parking lot). A way to ensure that there will be no accidents can be to require that every car will drive in a straight line, while if a change of heading occurs, it must be done when there are no close cars in my surrounding. The rationale behind this is that a car can
 30 predict what other cars will do, and behave accordingly. If other cars deviate from this prediction (by changing heading), it is done with a long enough distance and therefore there may be enough time to correct the prediction. When there is lane structure, it may enable smarter predictions on what other cars will do. If there is no lane structure at all, a car will continue according to its current heading. Technically speaking, this is equivalent to assigning every car to a virtual straight route according to its
 35 heading. Next, consider the scenario in a large unstructured roundabout (e.g., around the Arc de Triomphe in Paris). Here, a sensible prediction is to assume that a car will continue according to the geometry of the roundabout, while keeping its offset. Technically, this is equivalent to assigning every car to a virtual arc route according to its current offset from the center of the roundabout.

[0591] The above described driving policy system (e.g., the RL system) may be implemented together with one or more of the described accident liability rules to provide a navigational system that takes into account potential accident liability when deciding on a particular navigational instruction to implement. Such rules may be applied during the planning phase; e.g., within a set of programmed instructions or within a trained model such that a proposed navigational action is developed by the system already in compliance with the rules. For example, a driving policy module may account for or be trained with, for example, one or more navigational rules upon which RSS is based. Additionally or alternatively, the RSS safety constraint may be applied as a filter layer through which all proposed navigational actions proposed by the planning phase are tested against the relevant accident liability rules to ensure that the proposed navigational actions are in compliance. If a particular action is in compliance with the RSS safety constraint, it may be implemented. Otherwise, if the proposed navigational action is not in compliance with the RSS safety constraint (e.g., if the proposed action could result in accident liability to the host vehicle based on one or more of the above-described rules), then the action is not taken.

[0592] In practice, a particular implementation may include a navigation system for a host vehicle. The host vehicle may be equipped with an image capture device (e.g., one or more cameras such as any of those described above) that, during operation, captures images representative of an environment of the host vehicle. Using the image information, a driving policy may take in a plurality of inputs and output a planned navigational action for accomplishing a navigational goal of the host vehicle. The driving policy may include a set of programmed instructions, a trained network, etc., that may receive various inputs (e.g., images from one or more cameras showing the surroundings of the host vehicle, including target vehicles, roads, objects, pedestrians, etc.; output from LIDAR or RADAR systems; outputs from speed sensors, suspension sensors, etc.; information representing one or more goals of the host vehicle--e.g., a navigational plan for delivering a passenger to a particular location, etc.). Based on the input, the processor may identify a target vehicle in the environment of the host vehicle, e.g., by analyzing camera images, LIDAR output, RADAR output, etc. In some embodiments, the processor may identify a target vehicle in the environment of the host vehicle by analyzing one or more inputs, such as one or more camera images, LIDAR output, and/or RADAR output. Further, in some embodiments, the processor may identify a target vehicle in the environment of the host vehicle based on an agreement of a majority or combination of sensor inputs (e.g., by analyzing one or more camera images, LIDAR output, and/or RADAR output, and receiving a detection result identifying the target vehicle based on a majority agreement or combination of the inputs).

[0593] Based on the information available to the driving policy module, an output may be provided in the form of one or more planned navigational actions for accomplishing a navigational goal of the host vehicle. In some embodiments, the RSS safety constraint may be applied as a filter of the planned navigational actions. That is, the planned navigational action, once developed, can be tested against at least one accident liability rule (e.g., any of the accident liability rules discussed above) for determining potential accident liability for the host vehicle relative to the identified target vehicle. And,

as noted, if the test of the planned navigational action against the at least one accident liability rule indicates that potential accident liability may exist for the host vehicle if the planned navigational action is taken, then the processor may cause the host vehicle not to implement the planned navigational action. On the other hand, if the test of the planned navigational action against the at least one accident liability rule indicates that no accident liability would result for the host vehicle if the planned navigational action is taken, then the processor may cause the host vehicle to implement the planned navigational action.

[0594] In some embodiments, the system may test a plurality of potential navigational actions against the at least one accident liability rule. Based on the results of the test, the system may filter the potential navigational actions to a subset of the plurality of potential navigational actions. For example, in some embodiments, the subset may include only the potential navigational actions for which the test against the at least one accident liability rule indicates that no accident liability would result for the host vehicle if the potential navigational actions were taken. The system may then score and/or prioritize the potential navigational actions without accident liability and select one of the navigational actions to implement based on, for example, an optimized score, or a highest priority. The score and/or priority may be based, for example, one or more factors, such as the potential navigational action viewed as being the safest, most efficient, the most comfortable to passengers, etc.

[0595] In some instances, the determination of whether to implement a particular planned navigational action may also depend on whether a default emergency procedure would be available in a next state following the planned action. If a DEP is available, the RSS filter may approve the planned action. On the other hand, if a DEP would not be available, the next state may be deemed an unsafe one, and the planned navigational action may be rejected. In some embodiments, the planned navigational action may include at least one default emergency procedure.

[0596] One benefit of the described system is that to ensure safe actions by the vehicle, only the host vehicle's actions relative to a particular target vehicle need be considered. Thus, where more than one target vehicle is present, the planned action for the host vehicle may be tested for an accident liability rule sequentially with respect to the target vehicles in an influence zone in the vicinity of the host vehicle (e.g., within 25 meters, 50 meters, 100 meters, 200 meters, etc.). In practice, the at least one processor may be further programmed to: identify, based on analysis of the at least one image representative of an environment of the host vehicle (or based on LIDAR or RADAR information, etc.), a plurality of other target vehicles in the environment of the host vehicle and repeat the test of the planned navigational action against at least one accident liability rule for determining potential accident liability for the host vehicle relative to each of the plurality of other target vehicles. If the repeated tests of the planned navigational action against the at least one accident liability rule indicate that potential accident liability may exist for the host vehicle if the planned navigational action is taken, then the processor may cause the host vehicle not to implement the planned navigational action. If the repeated tests of the planned navigational action against the at least one accident liability rule indicate that no accident liability would result for the host vehicle if the planned navigational action is taken, then the processor may cause the host vehicle to implement the planned navigational action.

[0597] As noted, any of the rules described above can be used as the basis for the RSS safety test. In some embodiments, the at least one accident liability rule includes a following rule defining a distance behind the identified target vehicle within which the host vehicle may not proceed without a potential for accident liability. In other cases, the at least one accident liability rule includes a leading rule defining a distance forward of the identified target vehicle within which the host vehicle may not proceed without a potential for accident liability.

[0598] While the system described above can apply the RSS safety test to a single planned navigational action to test compliance with the rule that the host vehicle should not take any action for which it would be liable for a resulting accident, the test may be applied to more than one planned navigational action. For example, in some embodiments, the at least one processor, based on application of at least one driving policy may determine two or more planned navigational actions for accomplishing a navigational goal of the host vehicle. In these situations, the processor may test each of the two or more planned navigational actions against at least one accident liability rule for determining potential accident liability. And, for each of the two or more planned navigational actions, if the test indicates that potential accident liability may exist for the host vehicle if a particular one of the two or more planned navigational actions is taken, the processor may cause the host vehicle not to implement the particular one of the planned navigational actions. On the other hand, for each of the two or more planned navigational actions, if the test indicates that no accident liability would result for the host vehicle if a particular one of the two or more planned navigational actions is taken, then the processor may identify the particular one of the two or more planned navigational actions as a viable candidate for implementation. Next, the processor may select a navigational action to be taken from among the viable candidates for implementation based on at least one cost function and cause the host vehicle to implement the selected navigational action.

[0599] Because an implementation of RSS relates to a determination of relative potential liability for accidents between the host vehicle and one or more target vehicles, along with testing planned navigational actions for safety compliance, the system may track accident liability potential for encountered vehicles. For example, not only may the system be able to avoid taking an action for which a resulting accident would result in liability to the host vehicle, but the host vehicle systems may also be able to track one or more target vehicles and identify and track which accident liability rules have been broken by those target vehicles. In some embodiments, an accident liability tracking system for a host vehicle may include at least one processing device programmed to receive, from an image capture device, at least one image representative of an environment of the host vehicle and analyze the at least one image to identify a target vehicle in the environment of the host vehicle. Based on analysis of the at least one image, the processor may include programming to determine one or more characteristics of a navigational state of the identified target vehicle. The navigational state may include various operational characteristics of the target vehicle, such as vehicle speed, proximity to a center of a lane, lateral velocity, direction of travel, distance from the host vehicle, heading, or any other parameter that may be used to determine potential accident liability based on any of the rules described above. The processor may

compare the determined one or more characteristics of the navigational state of the identified target vehicle to at least one accident liability rule (e.g., any of the rules described above, such as winning by lateral velocity, directional priority, winning by proximity to lane center, following or leading distance and cut-in, etc.). Based on comparison of the state to one or more rules, the processor may store at least one value indicative of potential accident liability on the part of the identified target vehicle. And in the case of an accident, the processor may provide an output of the stored at least one value (e.g., via any suitable data interface, either wired or wireless). Such an output may be provided, for example, after an accident between the host vehicle and at least one target vehicle, and the output may be used for or may otherwise provide an indication of liability for the accident.

[0600] The at least one value indicative of potential accident liability may be stored at any suitable time and under any suitable conditions. In some embodiments, the at least one processing device may assign and store a collision liability value for the identified target vehicle if it is determined that the host vehicle cannot avoid a collision with the identified target vehicle.

[0601] The accident liability tracking capability is not limited to a single target vehicle, but rather can be used to track potential accident liability for a plurality of encountered target vehicles. For example, the at least one processing device may be programmed to detect a plurality of target vehicles in the environment of the host vehicle, determine navigational state characteristics for each of the plurality of target vehicles, and determine and store values indicative of potential accident liability on the part of respective ones of the plurality of target vehicles based on comparisons of the respective navigational state characteristics for each of the target vehicles to the at least one accident liability rule. As noted, the accident liability rules used as the basis for liability tracking may include any of the rules described above or any other suitable rule. For example, the at least one accident liability rule may include a lateral velocity rule, a lateral position rule, a driving direction priority rule, a traffic light-based rule, a traffic sign-based rule, a route priority rule, etc. The accident liability tracking function may also be coupled with safe navigation based on RSS considerations (e.g., whether any action of the host vehicle would result in potential liability for a resulting accident).

[0602] In addition to navigating based on accident liability considerations according to RSS, navigation can also be considered in terms of vehicle navigational states and a determination of whether a particular, future navigational state is deemed safe (e.g., whether a DEP exists such that accidents may be avoided or any resulting accident will not be deemed the fault of the host vehicle, as described in detail above). The host vehicle can be controlled to navigate from safe state to safe state. For example, in any particular state, the driving policy may be used to generate one or more planned navigational actions, and those actions may be tested by determining if the predicted future states corresponding to each planned action would offer a DEP. If so, the planned navigational action or actions providing the DEP may be deemed safe and may qualify for implementation.

[0603] In some embodiments, a navigation system for a host vehicle may include at least one processing device programmed to: receive, from an image capture device, at least one image representative of an environment of the host vehicle; determine, based on at least one driving policy, a

planned navigational action for accomplishing a navigational goal of the host vehicle; analyze the at least one image to identify a target vehicle in the environment of the host vehicle; test the planned navigational action against at least one accident liability rule for determining potential accident liability for the host vehicle relative to the identified target vehicle; if the test of the planned navigational action against the at least one accident liability rule indicates that potential accident liability may exist for the host vehicle if the planned navigational action is taken, then cause the host vehicle not to implement the planned navigational action; and if the test of the planned navigational action against the at least one accident liability rule indicates that no accident liability would result for the host vehicle if the planned navigational action is taken, then cause the host vehicle to implement the planned navigational action.

[0604] In some embodiments, a navigation system for a host vehicle, may include at least one processing device programmed to: receive, from an image capture device, at least one image representative of an environment of the host vehicle; determine, based on at least one driving policy, a plurality of potential navigational actions for the host vehicle; analyze the at least one image to identify a target vehicle in the environment of the host vehicle; test the plurality of potential navigational actions against at least one accident liability rule for determining potential accident liability for the host vehicle relative to the identified target vehicle; select one of the potential navigational actions for which the test indicates that no accident liability would result for the host vehicle if the selected potential navigational action is taken; and cause the host vehicle to implement the selected potential navigational action. In some instances, the selected potential navigational action may be selected from a subset of the plurality of potential navigational actions for which the test indicates that no accident liability would result for the host vehicle if any of the subset of the plurality of potential navigational action were taken. Further, in some instances, the selected potential navigational action may be selected according to a scoring parameter.

[0605] In some embodiments, a system for navigating a host vehicle may include at least one processing device programmed to receive, from an image capture device, at least one image representative of an environment of the host vehicle and determine, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle. The processor may also analyze the at least one image to identify a target vehicle in the environment of the host vehicle; determine a next-state distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken; determine a current maximum braking capability of the host vehicle and a current speed of the host vehicle; determine a current speed of the target vehicle and assume a maximum braking capability of the target vehicle based on at least one recognized characteristic of the target vehicle; and implement the planned navigational action if, given the maximum braking capability of the host vehicle and current speed of the host vehicle, the host vehicle can be stopped within a stopping distance that is less than the determined next-state distance summed together with a target vehicle travel distance determined based on the current speed of the target vehicle and the assumed maximum braking capability of the target vehicle. The stopping distance may further include a distance over which the host vehicle travels during a reaction time without braking.

[0606] The recognized characteristic of the target vehicle upon which the maximum braking capability of the target vehicle is determined may include any suitable characteristic. In some embodiments, the characteristic may include a vehicle type (e.g., motorcycle, car, bus, truck, each of which may be associated with different braking profiles), vehicle size, a predicted or known vehicle weight, a vehicle model (e.g., that may be used to look up a known braking capability), etc.

[0607] In some cases, the safe state determination may be made relative to more than one target vehicle. For example, in some cases a safe state determination (based on distance and braking capabilities) may be based on two or more identified target vehicles leading a host vehicle. Such a determination may be useful especially where information regarding what is ahead of the foremost target vehicle is not available. In such cases, it may be assumed for purposes of determining a safe state, safe distance, and/or available DEP that the foremost detectable vehicle will experience an imminent collision with an immovable or nearly immovable obstacle, such that the target vehicle following may reach a stop more quickly than its own braking profile allows (e.g., the second target vehicle may collide with a first, foremost vehicle and therefore reach a stop more quickly than expected max braking conditions). In such cases, it may be important to base the safe state, safe following distance, DEP determination upon the location of the foremost identified target vehicle relative to the host vehicle.

[0608] In some embodiments, such a safe state to safe state navigation system may include at least one processing device programmed to receive, from an image capture device, at least one image representative of an environment of the host vehicle. Here, as with other embodiments, the image information captured by an image capture device (e.g., a camera) may be supplemented with information obtained from one or more other sensors, such as a LIDAR or RADAR system. In some embodiments, the image information used to navigate may even originate from a LIDAR or RADAR system, rather than from an optical camera. The at least one processor may determine, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle. The processor may analyze the at least one image (e.g., obtained from any of a camera, a RADAR, a LIDAR, or any other device from which an image of the environment of the host vehicle may be obtained, whether optically based, distance map-based, etc.) to identify a first target vehicle ahead of the host vehicle and a second target vehicle ahead of the first target vehicle. The processor may then determine a next-state distance between the host vehicle and the second target vehicle that would result if the planned navigational action was taken. Next, the processor may determine a current maximum braking capability of the host vehicle and a current speed of the host vehicle. The processor may implement the planned navigational action if, given the maximum braking capability of the host vehicle and the current speed of the host vehicle, the host vehicle can be stopped within a stopping distance that is less than the determined next-state distance between the host vehicle and the second target vehicle.

[0609] That is, if the host vehicle processor determines that there is enough distance to stop in a next-state distance between the leading visible target vehicle and the host vehicle, without collision or without collision for which responsibility would attach to the host vehicle and assuming the leading visible target vehicle will suddenly at any moment come to a complete stop, then the processor of the host

vehicle may take the planned navigational action. On the other hand, if there would be insufficient room to stop the host vehicle without collision, then the planned navigational action may not be taken.

[0610] Additionally, while the next-state distance may be used as a benchmark in some embodiments, in other cases, a different distance value may be used to determine whether to take the planned navigational action. In some cases, as in the one described above, the actual distance in which the host vehicle may need to be stopped to avoid a collision may be less than the predicted next-state distance. For example, where the leading, visible target vehicle is followed by one or more other vehicles (e.g., the first target vehicle in the example above), the actual predicted required stopping distance would be the predicted next-state distance less the length of the target vehicle(s) following the leading visible target vehicle. If the leading visible target vehicle comes to an immediate stop, it may be assumed that the following target vehicles would collide with the leading visible target vehicle and, therefore, they too would need to be avoided by the host vehicle to avoid a collision. Thus, the host vehicle processor can evaluate the next-state distance less the summed lengths of any intervening target vehicles between the host vehicle and the leading, visible/detected target vehicle to determine whether there would be sufficient space to bring the host vehicle to a halt under max braking conditions without a collision.

[0611] In other embodiments, the benchmark distance for evaluating a collision between the host vehicle and one or more leading target vehicles may be greater than the predicted next-state distance. For example, in some cases, the leading visible/detected target vehicle may come to a quick, but not immediate stop, such that the leading visible/detected target vehicle travels a short distance after the assumed collision. For example, if that vehicle hits a parked car, the colliding vehicle may still travel some distance before coming to a complete stop. The distance traveled after the assumed collision may be less than an assumed or determined minimum stopping distance for the relevant target vehicle. Thus, in some cases, the processor of the host vehicle may lengthen the next-state distance in its evaluation of whether to take the planned navigational action. For example, in this determination, the next-state distance may be increased by 5%, 10%, 20%, etc. or may be supplemented with a predetermined fixed distance (10 m, 20 m, 50 m, etc.) to account for a reasonable distance that the leading/visible target vehicle may travel after an assumed imminent collision.

[0612] In addition to lengthening the next-state distance in the evaluation by an assumed distance value, the next-state distance may be modified by both accounting for a distance traveled after collision by the leading visible/detected target vehicle and the lengths of any target vehicles following the leading visible/detected target vehicle (which may be assumed to pile up with the leading visible/detected vehicle after its sudden stop).

[0613] In addition to basing the determination on whether to take the planned navigational action on the next-state distance between the host vehicle and the leading visible/detected target vehicle (as modified by considering the post-collision movement of the leading visible/detected target vehicle and/or the lengths of vehicles following the leading visible/detected target vehicle), the host vehicle may continue to account for the braking capability of one or more leading vehicles in its determination. For example, the host vehicle processor may continue to determine a next-state distance between the host

vehicle and the first target vehicle (e.g., a target vehicle following the leading visible/detected target vehicle) that would result if the planned navigational action was taken; determine a current speed of the first target vehicle and assume a maximum braking capability of the first target vehicle based on at least one recognized characteristic of the first target vehicle; and not implement the planned navigational action if, given the maximum braking capability of the host vehicle and the current speed of the host vehicle, the host vehicle cannot be stopped within a stopping distance that is less than the determined next-state distance between the host vehicle and the first target vehicle summed together with a first target vehicle travel distance determined based on the current speed of the first target vehicle and the assumed maximum braking capability of the first target vehicle. Here, as in the examples described above, the recognized characteristic of the first target vehicle may include a vehicle type, a vehicle size, a vehicle model, etc.

[0614] In some cases (e.g., through actions of other vehicles), the host vehicle may determine that a collision is imminent and unavoidable. In such cases, the processor of the host vehicle may be configured to select a navigational action (if available) for which the resulting collision would result in no liability to the host vehicle. Additionally or alternatively, the processor of the host vehicle may be configured to select a navigational action that would offer less potential damage to the host vehicle or less potential damage to a target object than the current trajectory or relative to one or more other navigational options. Further, in some cases, the host vehicle processor may select a navigational action based on considerations of the type of object or objects for which a collision is expected. For example, when faced with a collision with a parked car for a first navigational action or with an immovable object for a second navigational action, the action offering the lower potential damage to the host vehicle (e.g., the action resulting in a collision with the parked car) may be selected. When faced with a collision with a car moving in a similar direction as the host vehicle for a first navigational action or with a parked car for a second navigational action, the action offering the lower potential damage to the host vehicle (e.g., the action resulting in a collision with the moving car) may be selected. When faced with a collision with a pedestrian as a result of a first navigational action or with any other object for a second navigational action, the action offering any alternative to colliding with a pedestrian may be selected.

[0615] In practice, the system for navigating a host vehicle may include at least one processing device programmed to receive, from an image capture device, at least one image representative of an environment of the host vehicle (e.g., a visible image, LIDAR image, RADAR image, etc.); receive from at least one sensor an indicator of a current navigational state of the host vehicle; and determine, based on analysis of the at least one image and based on the indicator of the current navigational state of the host vehicle, that a collision between the host vehicle and one or more objects is unavoidable. The processor may evaluate available alternatives. For example, the processor may determine, based on at least one driving policy, a first planned navigational action for the host vehicle involving an expected collision with a first object and a second planned navigational action for the host vehicle involving an expected collision with a second object. The first and second planned navigational actions may be tested against at least one accident liability rule for determining potential accident liability. If the test of the first planned

navigational action against the at least one accident liability rule indicates that potential accident liability may exist for the host vehicle if the first planned navigational action is taken, then the processor may cause the host vehicle not to implement the first planned navigational action. If the test of the second planned navigational action against the at least one accident liability rule indicates that no accident liability would result for the host vehicle if the second planned navigational action is taken, then the processor may cause the host vehicle to implement the second planned navigational action. The objects may include other vehicles or non-vehicle objects (e.g., road debris, trees, poles, signs, pedestrians, etc.).

[0616] The following figures and discussion provide examples of various scenarios that may occur when navigating and implementing the disclosed systems and methods. In these examples, a host vehicle may avoid taking an action that would result in blame attributable to the host vehicle for a resulting accident if the action is taken.

[0617] Figs. 28A and 28B illustrate example following scenarios and rules. As shown in Fig. 28A, the region surrounding vehicle 2804 (e.g., a target vehicle) represents a minimum safe distance corridor for vehicle 2802 (e.g., a host vehicle), which is traveling in the lane a distance behind vehicle 2804. According to one rule consistent with disclosed embodiments, to avoid an accident in which blame is attributable to vehicle 2802, vehicle 2802 must maintain a minimum safe distance by remaining in the region surrounding vehicle 2802. In contrast, as shown in Fig. 28B, if vehicle 2804 brakes, then vehicle 2802 will be at fault if there is an accident.

[0618] Figs. 29A and 29B illustrate example blame in cut-in scenarios. In these scenarios, safe corridors around vehicle 2902 determine the fault in cut-in maneuvers. As shown in Fig. 29A, vehicle 2902 is cutting in front of vehicle 2902, violating the safe distance (depicted by the region surrounding vehicle 2902) and therefore is at fault. As shown in Fig. 29B, vehicle 2902 is cutting in front of vehicle 2902, but maintains a safe distance in front of vehicle 2904.

[0619] Figs. 30A and 30B illustrate example blame in cut-in scenarios. In these scenarios, safe corridors around vehicle 3004 determine whether vehicle 3002 is at fault. In Fig. 30A, vehicle 3002 is traveling behind vehicle 3006 and changes into the lane in which target vehicle 3004 is traveling. In this scenario, vehicle 3002 violates a safe distance and therefore is at fault if there is an accident. In Fig. 30B, vehicle 3002 cuts in behind vehicle 3004 and maintains a safe distance.

[0620] Figs. 31A-31D illustrate example blame in drifting scenarios. In Fig. 31A, the scenario starts with a slight lateral maneuver by vehicle 3104, cutting-in to the wide corridor of vehicle 3102. In Fig. 31B, vehicle 3104 continues cutting into the normal corridor of the vehicle 3102, violating a safe distance region. Vehicle 3104 is to blame if there is an accident. In Fig. 31C, vehicle 3104 maintains its initial position, while vehicle 3102 moves laterally "forcing" a violation of a normal safe distance corridor. Vehicle 3102 is to blame if there is an accident. In Fig. 31B, vehicle 3102 and 3104 move laterally towards each other. The blame is shared by both vehicles if there is an accident.

[0621] Figs. 32A and 32B illustrate example blame in two-way traffic scenarios. In Fig. 32A, vehicle 3202 overtakes vehicle 3206, and vehicle 3202 has performed a cut-in maneuver maintaining a safe distance from vehicle 3204. If there is an accident, vehicle 3204 is to blame for not braking with

reasonable force. In Fig. 32B, vehicle 3202 cuts-in without keeping safe longitudinal distance from vehicle 3204. In case of an accident, vehicle 3202 is to blame.

[0622] Figs. 33A and 33B illustrate example blame in two-way traffic scenarios. In Fig. 33A, vehicle 3302 drifts into the path of oncoming vehicle 3204, maintaining a safe distance. In case of an accident, vehicle 3204 is to blame for not braking with reasonable force. In Fig. 33B, vehicle 3202 drifts into the path of the oncoming vehicle 3204, violating a safe longitudinal distance. In case of an accident, vehicle 3204 is to blame.

[0623] Figs. 34A and 34B illustrate example blame in route priority scenarios. In Fig. 34A, vehicle 3202 runs a stop sign. Blame is attributed to vehicle 3202 for not respecting the priority assigned to vehicle 3204 by the traffic light. In Fig. 34B, although vehicle 3202 did not have priority, it was already in the intersection when vehicle 3204's light turned green. If vehicle 3204 hits 3202, vehicle 3204 would be to blame.

[0624] Figs. 35A and 35B illustrate example blame in route priority scenarios. In Fig. 35A, vehicle 3502 backing-up into the path of an oncoming vehicle 3504. Vehicle 3502 performs a cut-in maneuver maintaining a safe distance. In case of an accident, vehicle 3504 is to blame for not braking with reasonable force. In Fig. 35B, vehicle 3502 car cuts-in without keeping a safe longitudinal distance. In case of an accident, vehicle 3502 is to blame.

[0625] Figs. 36A and 36B illustrate example blame in route priority scenarios. In Fig. 36A, vehicle 3602 and vehicle 3604 are driving in the same direction, while vehicle 3602 turns left across the path of vehicle 3604. Vehicle 3602 performs cut-in maneuver maintaining safe distance. In case of an accident, vehicle 3604 is to blame for not braking with reasonable force. In Fig. 36B, vehicle 3602 cuts-in without keeping safe longitudinal distance. In case of an accident, vehicle 3602 is to blame.

[0626] Figs. 37A and 37B illustrate example blame in route priority scenarios. In Fig. 37A, vehicle 3702 wants to turn left, but must give way to the oncoming vehicle 3704. Vehicle 3702 turns left, violating safe distance with respect to vehicle 3704. Blame is on vehicle 3702. In Fig. 37B, vehicle 3702 turns left, maintaining a safe distance with respect to vehicle 3704. In case of an accident, vehicle 3704 is to blame for not braking with reasonable force.

[0627] Figs. 38A and 38B illustrate example blame in route priority scenarios. In Fig. 38A, vehicle 3802 and vehicle 3804 are driving straight, and vehicle 3802 has a stop sign. Vehicle 3802 enters the intersection, violating a safe distance with respect to vehicle 3804. Blame is on vehicle 3802. In Fig. 38B, vehicle 3802 enters the intersection while maintaining a safe distance with respect to vehicle 3804. In case of an accident, vehicle 3804 is to blame for not braking with reasonable force.

[0628] Figs. 39A and 39B illustrate example blame in route priority scenarios. In Fig. 39A, vehicle 3902 wants to turn left, but must give way to vehicle 3904 coming from its right. Vehicle 3902 enters the intersection, violating the right-of-way and a safe distance with respect to vehicle 3904. Blame is on vehicle 3902. In Fig. 39B, vehicle 3902 enters the intersection while maintaining the right-of-way and a safe distance with respect to vehicle 3904. In case of an accident, vehicle 3904 is to blame for not braking with reasonable force.

[0629] Figs. 40A and 40B illustrate example blame in traffic light scenarios. In Fig. 40A, vehicle 4002 is running a red light. Blame is attributed to vehicle 4002 for not respecting the priority assigned to vehicle 4004 by the traffic light. In Fig. 40B, although vehicle 4002 did not have priority, it was already in the intersection when the light for vehicle 4004 turned green. If vehicle 4004 hits vehicle 4002, vehicle 4004 would be to blame.

[0630] Figs. 41A and 41B illustrate example blame in traffic light scenarios. Vehicle 4102 is turning left across the path of the oncoming vehicle 4104. Vehicle 4104 has priority. In Fig. 41, vehicle 4102 turns left, violating a safe distance with respect to vehicle 4104. Blame is attributed to vehicle 4102. In Fig. 41B, vehicle 4102 turns left, maintaining a safe distance with respect to vehicle 4104. In case of an accident, vehicle 4104 is to blame for not braking with reasonable force.

[0631] Figs. 42A and 42B illustrate example blame in traffic light scenarios. In Fig. 42A, vehicle 4202 is turning right, cutting into the path of vehicle 4204, which is driving straight. Right-on-red is assumed to be a legal maneuver, but vehicle 4204 has right of way, as vehicle 4202 violates a safe distance with respect to vehicle 4204. Blame is attributed to vehicle 4202. In Fig. 42B, vehicle 4202 turns right, maintaining a safe distance with respect to vehicle 4204. In case of an accident, vehicle 4204 is to blame for not braking with reasonable force.

[0632] Figs. 43A-43C illustrate example vulnerable road users (VRUs) scenarios. Accidents with animals or VRUs where the vehicle performs a maneuver are treated as a variation of a cut-in, where the default blame is on the car, with some exceptions. In Fig. 43A, vehicle 4302 cuts into the path of an animal (or VRU) while maintaining safe distance and ensuring an accident can be avoided. In Fig. 43B, vehicle 4302 cuts into the path of an animal (or VRU) violating safe distance. Blame is attributed to vehicle 4302. In Fig. 43C, vehicle 4302 notices the animal and stops, giving the animal sufficient time to stop. If the animal hits the car, the animal is to blame.

[0633] Figs. 44A-44C illustrate example vulnerable road users (VRUs) scenarios. In Fig. 44A, vehicle 4402 is turning left at a signalized intersection and encounters a pedestrian in the crosswalk. Vehicle 4402 has a red light and the VRU has a green. Vehicle 4402 is at fault. In Fig. 44B, vehicle 4402 has a green light, and the VRU has a red light. If the VRU enters the crosswalk, the VRU is at fault. In Fig. 44C, vehicle 4402 has a green light, and the VRU has a red light. If the VRU was already in the crosswalk, vehicle 4402 is it fault.

[0634] Figs. 45A-45C illustrate example vulnerable road users (VRUs) scenarios. In Fig. 45A, vehicle 4402 is turning right and encounters a cyclist. The cyclist has a green light. Vehicle 4502 is at fault. In Fig. 45B, the cyclist has a red light. If the cyclist enters the intersection, the cyclist is at fault. In Fig. 45C, the cyclist has a red light, but was already in the intersection. Vehicle 4502 is at fault.

[0635] Figs. 46A-46D illustrate example vulnerable road users (VRUs) scenarios. Accidents with VRUs where a vehicle does not perform performance a maneuver are blamed on the car by default, with some exceptions. In Fig. 46A, vehicle 4602 must always make sure to maintain safe distance and ensuring an accident can be avoided with a VRU. In Fig. 46B, if vehicle 4602 does not maintain safe distance, vehicle 4602 is to blame. In Fig. 46C, if vehicle 4602 does not maintain sufficiently low speed

as to avoid colliding with a VRU that is potentially occluded by vehicle 5604, or drives above the legal limit, vehicle 4602 is to blame. In Fig. 46D, in another scenario with a potential occlusion of a VRU by vehicle 4604, if vehicle 4602 maintains sufficiently low speed, but the VRUs speed is above a reasonable threshold, the VRU is to blame.

5 [0636] As disclosed herein, RSS defines a framework for multi-agent scenarios. Accidents with static objects, road departures, loss of control or vehicle failure are blamed on the host. RSS defines cautious maneuvers, which will not allow accidents with other objects, unless other objects maneuvers dangerously into the path of the host (in which case they are to blame). In the case of a sure-collision where blame is on the target, the host will apply its brakes. The system may consider evasive steering
10 only if the maneuver is “cautious” (perceived not to cause another accident).

[0637] Non-collision incidents include accidents initiated by vehicle fires, potholes, falling objects, etc. In these cases, the blame may default to the host, except for scenarios which the host can avoid, such as potholes and falling objects, which may be classified as a “static object” scenario, assuming they become visible at safe distance or that a cautious evasive maneuver exists. In multi-agent
15 scenarios where the host vehicle was stationary, the host is not to blame. In this case the target essentially performed a non-safe cut-in. For example, if a cyclist rides into a stationary car, the host is not to blame.

[0638] RSS also includes guidelines for assigning blame where the road is not structured clearly, such as parking lots or wide roundabouts without lane marks. In these unstructured road scenarios, blame is assigned by examining deviations of each vehicle from its path to determine if they
20 allowed sufficient distance to allow other objects in the area to adjust.

[0639] Model for Safe and Scalable Self-Driving Vehicles

[0640] As previously discussed, a model called “Responsibility Sensitive Safety” (RSS) may formalize the notion of “accident blame.” RSS may be interpretable and explainable, and may incorporate a sense of “responsibility” into the actions of a robotic agent.

25 [0641] In some embodiments, RSS is constructed by formalizing the following four rules: (1) keep a safe distance from the car in front, such that if the car in front brakes abruptly you will be able to stop in time; (2) keep a safe distance from cars on either side, and, when performing lateral maneuvers and cutting-in to another car’s trajectory, leave the other car enough space to respond; (3) respect “right-of-way” rules, where “right-of-way” is given not taken; and (4) be cautious of occluded areas, for
30 example, an area behind a parked car.

[0642] The discussion below presents a formalization of the above rules. In some aspects, the formal model should satisfy soundness, e.g., when the model says that the self-driving car is not responsible for an accident, the model should match the “common sense” of human judgement, and usefulness, e.g., a policy that guarantees to never cause accidents while still maintaining normal flow of
35 traffic. In some embodiments, the model assigns responsibility on the self-driving car in fuzzy scenarios, possibly resulting in extra cautiousness, as long as the model is still useful.

[0643] Satisfying each of these requirements, soundness and uniqueness, individually is trivial — a model that always assign responsibility to the self-driving car is sound but not useful, while a model

that never assign responsibility to the self-driving car is useful but not sound. Embodiments of the RSS model satisfy both soundness and usefulness.

[0644] Safe Distance

[0645] The first basic concept formalized in the discussion of RSS is that the trailing car is always at fault if the trailing car hits a leading car. For example, two cars c_f , c_r , may be driving at the same speed, one behind the other, along a straight road, without performing any lateral maneuvers. Assume c_f , the car at the front, suddenly brakes because of an obstacle appearing on the road, and manages to avoid it. Unfortunately, c_r did not keep enough of a distance from c_f , is not able to respond in time, and crashes into c_f 's rear side. In this example, the blame is on c_r ; it is the responsibility of the rear car to keep safe distance from the front car, and to be ready for unexpected, yet reasonable, braking.

[0646] The following definition formalizes the concept of a "safe distance."

[0647] **Definition 27 (Safe longitudinal distance — same direction):** A longitudinal distance between a car c_r that drives behind another car c_f , where both cars are driving at the same direction, is safe with respect to a response time ρ if, for any braking of at most $a_{\max, \text{brake}}$ performed by c_f , c_r will accelerate by, at most, $a_{\max, \text{accel}}$ during the response time and then will brake by at least $a_{\min, \text{brake}}$ until a full stop then it will not collide with c_f .

[0648] In some embodiments, the safe longitudinal distance depends on parameters: ρ , $a_{\max, \text{accel}}$, $a_{\max, \text{brake}}$, $a_{\min, \text{brake}}$. These parameters may be determined, for example, by regulation. In some embodiments, the parameters are set differently for a robotic car and a human driven car. For example, the response time of a robotic car may be smaller than that of a human driver and a robotic car may brake more effectively than a typical human driver, hence $a_{\min, \text{brake}}$ may be a larger value for a robotic car than for a human-driven car. In other embodiments, the parameters may be set differently for different road conditions (wet road, ice, snow), which may be sensed, for example, by analysis of images acquired from one or more onboard cameras, or based on outputs from any other suitable sensors (e.g., windshield sensors, wheel slip sensors, etc.).

[0649] **Lemma 11** below calculates the safe distance as a function of the velocities of c_r , c_f and the parameters in **Definition 27**.

[0650] **Lemma 11:** Let c_r be a vehicle which is behind c_f on the longitudinal axis. Let ρ , $a_{\max, \text{brake}}$, $a_{\max, \text{accel}}$, $a_{\min, \text{brake}}$ be as in Definition 27. Let v_r , v_f be the longitudinal velocities of the cars. Then, the minimal safe longitudinal distance between the front-most point of c_r and the rear-most point of c_f is:

$$d_{\min} = \left[v_r \rho + \frac{1}{2} a_{\max, \text{accel}} \rho^2 + \frac{(v_r + \rho a_{\max, \text{accel}})^2}{2 a_{\max, \text{brake}}} - \frac{v_f^2}{2 a_{\max, \text{brake}}} \right]_+$$

[0651] **Proof:** Let d_0 denote the initial distance between c_r and c_f . Denote $v_{\rho, \max} = v_r + \rho a_{\max, \text{accel}}$. The velocity of the front car decreases with t at a rate $a_{\max, \text{brake}}$ (until arriving to zero or that a collision happens), while the velocity of the rear car increases in the time interval $[0, \rho]$ (until reaching $v_{\rho, \max}$) and then decreases to zero (e.g., a full stop or a collision) at a rate $a_{\min, \text{brake}} < a_{\max, \text{brake}}$. It follows

that if, at some point in time, the two cars have the same velocity, then from there on, the front car's velocity will be smaller, and the distance between them will be monotonically decreasing until both cars reach a full stop (where the "distance" can be negative if collision happens). Thus, the worst-case distance may happen either at time zero or when the two cars reach a full stop. In the former case, $d_0 >$

5 0. In the latter case, the distances the front and rear car will pass until a full stop may be $\frac{v_f^2}{a_{\max, \text{brake}}}$ and $v_r \rho + \frac{1}{2} a_{\max, \text{accel}} \rho^2 + \frac{v_{\rho, \max}^2}{2 a_{\min, \text{brake}}}$, respectively. At that point (when both cars reach a full stop), the distance between them should be larger than zero,

$$d_0 + \frac{v_f^2}{2 a_{\max, \text{brake}}} - \left(v_r \rho + \frac{1}{2} a_{\max, \text{accel}} \rho^2 + \frac{v_{\rho, \max}^2}{2 a_{\min, \text{brake}}} \right) > 0$$

[0652] Rearranging terms yields **Lemma 11**, or the minimum safe distance.

10 [0653] The preceding definition of a safe distance may be sound for the case that both the rear and front cars are driving at the same direction. Indeed, in this case, it is the responsibility of the rear car to keep a safe distance from the front car, and to be ready for unexpected, yet reasonable, braking. However, a definition is also needed for the scenario in which the two cars are driving in opposite directions.

15 [0654] Consider, for example, a car c_r , that is currently at a safe distance from a preceding car, c_f , that is stopped. Suddenly, c_f may reverse into a parking spot, causing a collision with c_r from behind. Here, even though c_r hit c_f from behind, the responsibility is not on the rear car c_r . To formalize this rule, the definitions of "rear and front" are defined so as not to apply to scenarios in which vehicles are moving toward each other (namely, when the signs of their longitudinal velocities are opposite). In
20 such cases, it may be expected that both cars decrease the absolute value of their velocities in order to avoid a crash.

[0655] Thus, the safe distance between cars that drive in opposite directions may be defined as the distance required so that, if both cars brake (after a given response time), then no crash will occur. However, it may be assumed that the car that drives in the opposite direction to the lane direction should
25 brake harder than the one driving in the correct direction. This leads to the following definition.

[0656] **Definition 28 (Safe longitudinal distance — opposite directions:** Consider cars c_1 , c_2 driving on a lane with longitudinal velocities v_1, v_2 , where $v_2 < 0$ and $v_1 \geq 0$ (the sign of the longitudinal velocity is according to the allowed direction of driving on the lane). The longitudinal distance between the cars is safe with respect to a response time ρ , braking parameters $a_{\min, \text{brake}}$,
30 $a_{\min, \text{brake}, \text{correct}}$, and an acceleration parameter $a_{\max, \text{accel}}$, if in case c_1, c_2 will increase the absolute value of their velocities at rate $a_{\max, \text{accel}}$ during the response time, and from there on will decrease the absolute value of their velocities at rate $a_{\min, \text{brake}, \text{correct}}, a_{\min, \text{brake}}$, respectively, until a full stop, then there will not be a collision.

[0657] A calculation of the safe distance for the case of opposite directions may be given in
35 the lemma below (whose proof is straightforward, and hence omitted).

[0658] **Lemma 12:** Consider the notation given in **Definition 28**. Define $v_{1,\rho} = v_1 + \rho a_{\max, \text{ accel}}$ and $v_{2,\rho} = |v_2| + \rho a_{\max, \text{ accel}}$. Then, the minimal safe longitudinal distance between c_1 and c_2 is:

$$d_{\min} = \frac{v_1 + v_{1,\rho}}{2} \rho + \frac{v_{a,\rho}^2}{2a_{\min, \text{ brake, correct}}} + \frac{|v_2| + v_{2,\rho}}{2} \rho + \frac{v_{2,\rho}^2}{2a_{\min, \text{ brake}}}$$

5 [0659] Before a collision between two cars, the cars must first be at a non-safe distance.

Intuitively, the idea of the safe distance definitions is that if both cars will respond “properly” to violations of the safe distance then there cannot be a collision. If one of them does not respond “properly” then that care is responsible for the accident. To formalize this, it is first important to define the moment just before the cars start to be at a non-safe distance.

10 [0660] **Definition 29 (Dangerous Longitudinal Situation and Blame Time):** Let time t be dangerous for cars c_1, c_2 if the distance between them at time t is non-safe (according to Definition 27 or Definition 28). Given a dangerous time t , its Blame Time, denoted t_b , is the earliest non-dangerous time such that all the times in the interval $(t_b, t]$ are dangerous. For example, an accident can only happen at time t if it is dangerous, and thus the blame time of the accident is the blame time of t .

15 [0661] **Definition 30 (Proper response to dangerous longitudinal situations):** As above, let t be a dangerous time for cars c_1, c_2 and let t_b be the corresponding blame time. Assume the two cars comply with the following constraints on the longitudinal speed:

1. First, for example, if at the blame time, the two cars were driving at the same direction, and say that c_1 is the rear car, then

20 • c_1 acceleration must be at most $a_{\max, \text{ accel}}$ during the interval $[t_b, t_b + \rho)$ and at most $-a_{\min, \text{ brake}}$ from time $t_b + \rho$ until reaching a full stop, after which any non-positive acceleration is allowed. In other words, the following car is allowed to accelerate at a rate up to its maximum acceleration capability (e.g., full throttle level) during the reaction time period, but then, must decelerate at a rate associated with a braking rate between a selected minimum braking rate for the car and the car’s maximum braking rate. The selected minimum braking rate may be regulation dependent or selected based on any other suitable criteria.

25 • c_2 acceleration must be at least $-a_{\max, \text{ brake}}$ until reaching a full stop, after which, any non-negative acceleration is allowed. That is, the leading car is allowed to decelerate at a rate up to a deceleration rate associated with the leading car’s maximum braking capability.

30 2. Second, if at the blame time the two cars were driving at opposite directions, and, for example, c_2 was driving in the wrong direction (e.g., c_2 has a negative velocity), then:

35 • c_1 acceleration must be at most $a_{\max, \text{ accel}}$ during the interval $[t_b, t_b + \rho)$ and at most $-a_{\min, \text{ brake, correct}}$ from time $t_b + \rho$ until reaching a full stop, after which any non-positive acceleration is allowed. Car c_1 may accelerate toward car c_2 at

a maximum acceleration capability during the reaction time, but then, must decelerate at a rate associated with a braking rate between a selected minimum braking rate for the car and the car's maximum braking rate. The selected minimum braking rate may be regulation dependent or selected based on any other suitable criteria.

- c_2 acceleration must be at least $-a_{\max, \text{accel}}$ during the interval $[t_b, t_b + \rho)$ and at most $a_{\min, \text{brake}}$ from time $t_b + \rho$ until reaching a full stop, after which any non-negative acceleration is allowed. In other words, car c_2 may accelerate toward car c_1 at a rate up to a maximum acceleration rate for car c_2 during car c_2 's reaction time. After that, car c_2 must decelerate at a rate associated with a braking rate between a selected minimum braking rate for the car and the car's maximum braking rate. The selected minimum braking rate may be regulation dependent or selected based on any other suitable criteria. Additionally, the selected minimum braking rate of car c_2 may be different than the selected minimum braking rate of car c_1 (e.g., the minimum braking rate that must be applied by the car driving in the incorrect direction may be higher than the minimum braking rate required of the car driving in the correct direction).

[0662] As mentioned previously, assume that collisions can only happen at dangerous times. Thus, if once the distance between cars become non-safe they both apply their "proper response to dangerous situations" from the corresponding blame time (until the distance between them becomes safe again) then collision cannot happen. Hence the responsibility for accidents may be assigned to the agent(s) that did not respond properly to the dangerous situation preceding the collision. In some embodiments, the definition below may apply to scenarios in which no lateral maneuvers are performed (that is, each car keeps its lateral position in the lane).

[0663] In the next section, the definition may be extended to the case in which cars can perform lateral maneuvers.

[0664] **Definition 31 (Responsibility for an Accident — no lateral maneuvers):** Consider an accident between c_1 and c_2 at time t and let t_b be the corresponding blame time. Car c_1 is responsible for the accident if it did not follow the constraints defined by the "proper response to dangerous situations" (**Definition 30**) at some time in the interval $(t_b, t]$.

[0665] Additionally, it may be possible that both c_1 and c_2 share the responsibility for an accident (e.g., if both of them did not comply with the proper response constraints).

[0666] Responsibility in Lateral Maneuvers

[0667] In some embodiments, the RSS model may formally define responsibility when cars are performing lateral maneuvers. A first example assumes a straight road on a planar surface, where the lateral, longitudinal axes are the x, y axes, respectively. For example, as discussed above with reference to the lane-based coordinate system. In some embodiments, the longitudinal and lateral velocities may be the derivatives of the longitudinal and lateral positions on the straight virtual road obtained by a bijection

between the actual curved road and a straight road. Similarly, accelerations may be second derivatives of positions.

[0668] Unlike longitudinal velocity, which may be kept at a value of 0 for long periods of time (e.g., when the car is not moving), keeping lateral velocity at exact 0 may be difficult as cars usually
5 perform small lateral fluctuations. Thus, the model may introduce a robust notion of lateral velocity.

[0669] **Definition 32 (μ -lateral-velocity):** Consider a point located at a lateral location l at time t . Its μ -lateral velocity at time t is defined as follows. Let $t_{\text{out}} > t$ be the earliest future time at which the point's lateral position, denoted l_{out} , is either $l - \frac{\mu}{2}$ or $l + \frac{\mu}{2}$ (if no such time exists let $t_{\text{out}} = 4$).
10 If at some time $t' \in (t, t_{\text{out}})$ the point's lateral position is l , then the μ -lateral-velocity is 0. Otherwise, the μ -lateral-velocity is $\frac{l_{\text{out}} - l}{t_{\text{out}} - t}$.

[0670] In some embodiments, in order to have a collision between two vehicles, it is required that they are close in proximity both longitudinally and laterally. A safe longitudinal distance has been defined above.

[0671] **Definition 33 (Safe Lateral Distance):** The lateral distance between cars c_1, c_2
15 driving with lateral velocities v_1, v_2 is safe with respect to parameters $\rho, a_{\text{min, brake}}^{\text{lat}}, a_{\text{max, accel}}^{\text{lat}}, \mu$ if during the time interval $[0, \rho]$ the two cars will apply lateral acceleration of $a_{\text{max, accel}}^{\text{lat}}$ toward each other, and after that the two cars will apply lateral braking of $a_{\text{min, brake}}^{\text{lat}}$, until each car reaches zero lateral velocity, then the final lateral distance between them will be at least μ .

[0672] A calculation of the lateral safe distance may be given in the lemma below.

20 [0673] **Lemma 13:** Consider the notation given in **Definition 33**. Without loss of generality, assume that c_1 is to the left of c_2 . Define $v_{1,\rho} = v_1 + a_{\text{max, accel}}^{\text{lat}} \rho$ and $v_{2,\rho} = v_2 - a_{\text{max, accel}}^{\text{lat}} \rho$. Then, the minimal safe lateral distance between the right side of c_1 and the left part of c_2 is:

$$d_{\text{min}} = \mu + \left[\frac{v_1 + v_{1,\rho}}{2} \rho + \frac{v_{1,\rho}^2}{2a_{\text{min, brake}}^{\text{lat}}} - \left(\frac{v_2 + v_{2,\rho}}{2} \rho - \frac{v_{2,\rho}^2}{2a_{\text{min, brake}}^{\text{lat}}} \right) \right]_+$$

[0674] As previously mentioned, in order to have a collision between two cars, assume the
25 cars must be both at a non-safe longitudinal distance and at a non-safe lateral distance. Thus, the idea of the safe distance definitions is that if both cars respond "properly" to violations of safe distance then there cannot be a collision. If one car did not respond "properly" then that car is responsible for the accident. In this example, the model may use modified definitions of Dangerous Situation and Blame Time.

[0675] **Definition 34 (Dangerous Situation and Blame Time):** Let time t be dangerous for
30 cars c_1, c_2 if both the longitudinal and lateral distances between them are non-safe (according to Definition 27, Definition 28, and Definition 33). Given a dangerous time t , its Blame Time, denoted t_b , is the earliest non-dangerous time such that all the times in the interval $(t_b, t]$ are dangerous. In this example, an accident may only happen at time t if it is dangerous, and in that case the blame time of the accident is the blame time of t .

[0676] **Definition 35 (Proper response to dangerous situations):** Let t be a dangerous time for cars c_1 , c_2 and let t_b be the corresponding blame time. Assume the proper behavior of the two cars is to comply with the following constraints on the lateral/longitudinal speed:

- 5 1. If before the blame time there was a safe longitudinal distance between c_1 and c_2 then the longitudinal speed is constrained according to **Definition 30**.
2. If before the blame time there was a safe lateral distance between c_1 and c_2 , and without loss of generality, assuming that at that time c_1 was to the left of c_2 :
 - if $t \in [t_b, t_b + \rho)$ then both cars can do any lateral action as long as their lateral acceleration, a , satisfies $|a| \leq a_{\max, \text{ accel}}^{\text{lat}}$.
 - Otherwise, if $t > t_b + \rho$:
 - then before reaching μ -lateral-velocity of 0, c_1 must apply lateral acceleration of at most $-a_{\min, \text{ brake}}^{\text{lat}}$ and c_2 must apply lateral acceleration of at least $a_{\min, \text{ brake}}^{\text{lat}}$. In other words, c_1 must have at least a minimum lateral acceleration in one direction, and c_2 must have at least a minimum lateral acceleration in the opposite direction.
 - After reaching μ -lateral-velocity of 0, c_1 may have any non-positive μ -lateral-velocity and c_2 may have any non-negative μ -lateral-velocity.
3. Minimal evasive effort: in addition to the rules above, the following constraints may apply.
 - First, in the case of item (1) above: Suppose that at the blame time the two cars are driving at the same direction and c_2 is the car on the front. Then, after time $t_b + \rho$ and until reaching a μ -lateral-velocity of 0, it should brake laterally by at least $a_{\min, \text{ brake, evasive}}^{\text{lat}}$. And, after that it should stay at μ -lateral-velocity of 0.
 - In case of item (2) above: Suppose that c_1 responds properly until reaching a μ -lateral-velocity of 0, and during that time c_2 did not respond properly and is now in front of c_1 (namely, all points of c_2 are longitudinally in front of all points of c_1 , and some point of c_2 is at the same lateral position of some point of c_1). Let the first moment that this happens, the cut-in time, be denoted t_c . Then, if $t_b < t_c < t - \rho$ and currently c_1 's velocity is positive, then c_1 must brake longitudinally at a rate of at least $a_{\min, \text{ brake, evasive}}$.

[0677] In the above definition, all but the item titled "Minimal evasive effort" capture the essence of the assumptions in the definitions of safe distance, in the sense that if both cars respond properly then there will no collision. The "Minimal evasive effort" deals with cases in which extra caution is applied to prevent potential situations from arising in which responsibility might be shared. For example, when the rear car, c_1 , should keep a safe distance from the front car, c_2 , and c_2 is overly cautious and avoids lateral maneuvers so as to prevent collisions in case c_1 will not brake strong enough.

The second case is when, for example, c_2 performs a cut-in, at a non-safe longitudinal distance, while c_1 was behaving properly. In this case, c_2 did not behave correctly. Nevertheless, c_1 is expected to make an effort to avoid a potential collision, for example, by applying braking of at least $a_{\min, \text{brake, evasive}}$. This may reduce the duration of a dangerous situation (unless both cars are at a zero longitudinal and lateral velocity).

[0678] **Definition 36 (Responsibility for an Accident):** Consider, for example, an accident between c_1 and c_2 at time t and let t_b be the corresponding Blame Time. Car c may be responsible for the accident if it did not respond properly according to **Definition 35** at some time in $(t_b, t]$, as illustrated in Fig. 47A.

[0679] In Fig. 47A, the vertical lines around each car show the possible lateral position of the car if it will accelerate laterally during the response time and then will brake laterally. Similarly, the rectangles show the possible longitudinal positions of the car (e.g., if the car either brakes by $a_{\max, \text{brake}}$ or accelerates during the response time and then will brake by $a_{\min, \text{brake}}$). In the top two rows (rows a and b), before the blame time there was a safe lateral distance, hence the proper response is to brake laterally. The yellow car is already at μ -lateral-velocity of zero, hence only the red car brakes laterally. In the third row (row c), before the blame time there was a safe longitudinal distance, hence the proper response is for the car behind to brake longitudinally. Finally, in the fourth row (row d), before the blame time there was a safe longitudinal distance, in an oncoming scenario, hence both car should brake longitudinally.

[0680] The parameters $a_{\max, \text{accel}}^{\text{lat}}$, $a_{\max, \text{accel}}$, and $a_{\max, \text{brake}}$ do not necessarily reflect a physical limitation, but instead they may represent an upper bound on reasonable behavior expected from road users. In some embodiments, based on these definitions, if a driver does not comply with these parameters at a dangerous time he immediately becomes responsible for the accident.

[0681] For simplicity, assume that cars can immediately switch from applying a “lateral braking” of $a_{\min, \text{brake}}^{\text{lat}}$ to being at a μ -lateral-velocity of 0. This may not always be possible due to physical properties of the car. However, the important factor in the definition of proper response is that from time $t_b + \rho$ to the time the car reaches a μ -lateral-velocity of 0, the total lateral distance it will pass will not be larger than the one it would have passed had it applied a braking of $a_{\min, \text{brake}}^{\text{lat}}$ until a full stop. In some embodiments, achieving this goal by a real vehicle is possible by first braking at a stronger rate, and then decreasing lateral speed more gradually at the end. In some embodiments, this change in the braking rate may not affect the RSS model.

[0682] In some embodiments, the definitions above hold for vehicles of arbitrary shapes, by taking the worst-case with respect to all points of each car. For example, the model may account for semi-trailers or a car with an open door.

[0683] Multiple Geometry and Right-of-Way Rules

[0684] In some scenarios, there may be multiple different road geometries in one scene, or environment, that overlap in a certain area. Fig. 47B shows exemplary road geometries including a roundabout 4706, junctions 4704 and 4708, and merges onto highways 4702. In many such cases, one

route, shown in solid line, may have priority over others, shown in dashed line, and vehicles riding on that route have the right of way.

[0685] When facing scenarios of multiple route geometries, the above definitions may need to be adjusted. For example, consider the T-junction 4702 depicted in Fig. 47B, and suppose that there is a stop sign for the route denoted by the solid line. Suppose that c_1 is approaching the intersection on the dashed route and at the same time c_2 is approaching the intersection on the solid route. According to the dashed route's coordinate system, c_2 has a very large lateral velocity, hence c_1 might deduce that c_2 is already at a non-safe lateral distance, which implies that c_1 , driving on the prioritized route, must reduce speed in order to maintain a safe longitudinal distance to c_2 . This means that c_2 should be very conservative with respect to traffic that coming from the solid route. This is of course an unnatural behavior, as cars on the yellow route have the right-of-way in this case. Furthermore, even c_2 , who doesn't have the priority, should be able to merge into the junction as long as c_1 can stop in time (this will be crucial in dense traffic). This example shows that when c_1 drives on r_1 , it doesn't make sense to consider its position and velocity w.r.t the coordinate system of r_2 . As a result, the preceding concepts such as "what does it mean that c_1 is in front of c_2 ", and what does it mean to be at a non-safe distance may be generalized.

[0686] In the case of multiple geometries, the definition of safe lateral distance may be generalized. Applying the definition below to two routes of the same geometry may yield the same definition as in **Definition 33**. As referred to herein, a route may be a subset of \mathbb{R}^2 . Further, assume that two cars in the examples below, c_1, c_2 are driving on different routes, r_1, r_2 . In some situations (for example, the T-junction 4704 given in Fig. 47B), once there is exactly a single route r_1 such that both cars are assigned to the route, and the time is not dangerous, then from that moment on, the definitions are described solely with reference to r_1 .

[0687] **Definition 37 (Lateral Safe Distance for Two Routes of Different Geometry):**
 Consider, for example, vehicles c_1, c_2 driving on routes r_1, r_2 that intersect. For every $i \in \{1, 2\}$, let $[x_{i,\min}, x_{i,\max}]$ be the minimal and maximal lateral positions in r_i that c_i can be in, if during the time interval $[0, \rho)$ the car applies a lateral acceleration (with respect to r_i), $|a^{\text{lat}}| \leq a_{\max, \text{accel}}^{\text{lat}}$, and after that applies a lateral braking of at least $a_{\min, \text{brake}}^{\text{lat}}$ (again with respect to r_i), until reaching a zero lateral velocity (with respect to r_i). The lateral distance between c_1 and c_2 is safe if the restrictions of r_1, r_2 to the lateral intervals $[x_{1,\min}, x_{1,\max}], [x_{2,\min}, x_{2,\max}]$ are at a distance of at least μ . Where the restriction of r_i to the lateral intervals $[x_{i,\min}, x_{i,\max}]$ may be the subset of \mathbb{R}^2 obtained by all points $(x, y) \in r_i$ for which the semantic lateral position of (x, y) is in the interval $[x_{i,\min}, x_{i,\max}]$.

[0688] **Definition 38**, below, is an exemplary quantification of the ordering between two cars when no common longitudinal axis exists (e.g., if the cars are not driving one in front of the other).

[0689] **Definition 38 (Longitudinal Ordering for Two Routes of Different Geometry):**
 Consider, for example, $c_1, c_2, 4710$ and 4712 , respectively, driving on routes $r_1, r_2, 4714$ and 4716 , respectively, that intersect as shown in Fig. 47C. For example, c_1 (vehicle 4710) is longitudinally in front of c_2 (vehicle 4712) if either of the following holds: for every i , if both vehicles are on r_i then c_1 is in front

of c_2 according to r_1 ; or c_1 is outside r_2 and c_2 is outside r_1 , and the longitudinal distance from c_1 to the set $r_1 \cap r_2$, with respect to r_1 , is larger than the longitudinal distance from c_2 to the set $r_1 \cap r_2$, with respect to r_2 .

[0690] As a measure of robustness of the longitudinal ordering definition, for example, suppose that c_1, c_2 are at distances of 20, 20.1 meters, respectively, from the intersection. The definition holds because the definition is effectively being used only when there is a safe longitudinal distance between the two cars, and, thus, the ordering between the cars will be obvious. Furthermore, this situation is analogous to the non-robustness of ordering when two cars are driving side by side on a multi-lane highway road.

[0691] **Definition 39 (Longitudinal Safe Distance for Two Routes of Different Geometry):** Consider, for example, c_1, c_2 driving on routes r_1, r_2 that intersect. The longitudinal distance between c_1 and c_2 is safe if one of the following holds:

1. If for all $i \in \{1, 2\}$ such that r_i has no priority, if c_i will accelerate by $a_{\max, \text{accel}}$ for ρ seconds, and will then brake by $a_{\min, \text{brake}}$ until reaching zero longitudinal velocity (all with respect to r_i), then during this time c_i will remain outside of the other route.
2. Otherwise, if c_1 is in front of c_2 (according to **Definition 38**), then the cars are at a safe longitudinal distance. For example, if c_1 brakes by $a_{\max, \text{brake}}$ until reaching a zero velocity (with respect to r_1), and c_2 accelerates by at most $a_{\max, \text{accel}}$ for ρ seconds and then brakes by at least $a_{\min, \text{brake}}$ (with respect to r_2) until reaching a zero velocity, then c_1 will remain in front of c_2 (according to **Definition 38**).
3. Otherwise, consider a point $p \in r_1 \cap r_2$ such that, for $i \in \{1, 2\}$, the lateral position of p with respect to r_i is in $[x_{i, \min}, x_{i, \max}]$ (as defined in **Definition 37**). Let $[t_{i, \min}, t_{i, \max}]$ be all times such that c_i can arrive to the longitudinal position of p with respect to r_i if it will apply longitudinal accelerations in the range $[-a_{\max, \text{brake}}, a_{\max, \text{accel}}]$ during the first ρ seconds, and then will apply longitudinal braking in the range $[a_{\min, \text{brake}}, a_{\max, \text{brake}}]$ until reaching a zero velocity. Then, the vehicles may be at a safe longitudinal distance if for every such p , $[t_{1, \min}, t_{1, \max}]$ does not intersect $[t_{2, \min}, t_{2, \max}]$.

[0692] Fig. 47D is an exemplary illustration of **Definition 39**. Frame 4720 is an exemplary safe situation in which car 4722 has priority because car 4724 must stop prior to entering the roadway. Frame 4726 is another exemplary safe situation in which car 4722 is in front of car 4724, thus giving car 4724 time to stop if car 4722 stops or slows. In frame 4728, If car 4722 is at a full stop and car 4724 is at a full lateral stop, the situation is safe by the definition of item (3) of **Definition 39**.

[0693] **Definition 40 (Dangerous & Blame Times, Proper Response, and Responsibility for Routes of Different Geometry):** Consider, for example, vehicles c_1, c_2 driving on routes r_1, r_2 . Time t may be dangerous if both the lateral and longitudinal distances are non-safe (according to Definition 11 and Definition 13). The corresponding blame time may be the earliest non-dangerous time t_b such that all times in $(t_b, t]$ are dangerous. The proper response of c_1 and/or c_2 may depend on the situation immediately before the blame time:

- If the lateral distance was safe, then both cars should respond according to the description of lateral safe distance in **Definition 37**.
- Else, if the longitudinal distance was safe according to item (1) in **Definition 39**, then if a vehicle is on the prioritized route it can drive normally, and otherwise it must brake by at least $a_{\min, \text{brake}}$ if $t - t_b \geq \rho$.
- Else, if the longitudinal distance was safe according to item (2) in **Definition 39**, then c_1 can drive normally and c_2 must brake by at least $a_{\min, \text{brake}}$ if $t - t_b \geq \rho$.
- Else, if the longitudinal distance was safe according to item (3) in **Definition 39**, then both cars can drive normally if $t - t_b < \rho$, and otherwise, both cars should brake laterally and longitudinally by at least $a_{\min, \text{brake}}^{\text{lat}}$, $a_{\min, \text{brake}}$ (each one with respect to its own route).

[0694] Finally, if a collision does occur, then the responsibility may be on the vehicle(s) that did not comply with the proper response. In some embodiments, there are cases where the route used by another agent is unknown: for example, in Fig. 47E. In Fig. 47E, at the path of car 4730 at the intersection is unknown to car 4732. In such case, RSS may be obtained by checking all possibilities.

[0695] Traffic Lights

[0696] In some embodiments, a vehicle may encounter a situation involving a traffic light. It may not hold that “if one car’s route has the green light and the other car’s route has a red light, then the blame is on the one whose route has the red light.” Rather, consider for example the scenario depicted in Fig. 47F. Even if the car 4734’s route has a green light, it is not expected it to ignore car 4736 that is already in the intersection. The route that has a green light may have a priority over routes that have a red light. Therefore, the concept of a situation with a traffic light reduces to the route priority situation previously described. In some embodiments, the rule for route priority may be simplified conceptually as the right of way is given, not taken.

[0697] Unstructured Road

[0698] In some embodiments, a vehicle may encounter unstructured roads, for example, as shown in Fig. 47G. Consider first the scenario 4718. Here, while the partition of the road area to lanes is not well defined, the partition of the road to multiple routes (with a clear geometry for every route) is well defined. Since our definitions of responsibility only depend on the route geometry, they apply as is to such scenarios.

[0699] Next, consider the scenario where there is no route geometry at all (e.g. the parking lot 4738). Unlike the structured case, lateral and longitudinal directions are separated, this scenario requires two-dimensional trajectories.

[0700] **Definition 41 (Trajectories):** Consider, for example, a vehicle c riding on some road. A future trajectory of c may be represented as a function $\tau : \mathbb{R}_+ \rightarrow \mathbb{R}^2$, where $\tau(t)$ is the position of c in t seconds from the current time. The tangent vector to the trajectory at t , denoted $\tau'(t)$, is the Jacobian of τ at t . Denoting $t_s(\tau) = \sup\{t : \forall t_1 \in [0, t), \|\tau'(t_1)\| > 0\}$, namely, that t_s is the first time in which the vehicle will arrive to a full stop, where if no such t exists, set $t_s(\tau) = \infty$.

[0701] Dangerous situations may be identified based on the possibility of a collision between two trajectories, as described with reference to **Definition 42**.

[0702] **Definition 42 (Trajectory Collision):** Let τ_1, τ_2 be two future trajectories of c_1, c_2 , with corresponding stopping times $t_1 = t_s(\tau_1), t_2 = t_s(\tau_2)$. Given parameters ϵ, θ , then τ_1 and τ_2 do not collide, which may be denoted by $\tau_1 \cap \tau_2 = \emptyset$, if either of the following holds:

1. For every $t \in [0, \max(t_1, t_2)]$ then $\|\tau_1(t) - \tau_2(t)\| > \epsilon$.
2. For every $t \in [0, t_1]$ then $\|\tau_1(t) - \tau_2(t)\| > \epsilon$ and the absolute value of the angle between the vectors $(\tau_2(t_1) - \tau_1(t_1))$ and $\tau_2'(t_1)$ is at most θ .

[0703] Given a set of trajectories for c_1 , denoted T_1 , and a set of trajectories for c_2 , denoted T_2 , then $T_1 \cap T_2 = \emptyset$ if for every $(\tau_1, \tau_2) \in T_1 \times T_2$ then $\tau_1 \cap \tau_2 = \emptyset$.

[0704] The first item states, for example, that both vehicles will be away from each other until they are both at a full stop. The second item states, for example, that the vehicles will be away from each other until the first one is at a full stop, and at that time, the velocity vector of the second one points away from the first vehicle.

[0705] In some embodiments, the collision operator may not be commutative. For example, assume two cars are currently driving on a large circle in the same direction, where c_1 is closely behind c_2 , and consider τ_1 to be the trajectory in which c_1 brakes strongly and τ_2 is the trajectory in which c_2 continues at the same speed forever. Then, $\tau_1 \cap \tau_2 = \emptyset$ while $\tau_2 \cap \tau_1 \neq \emptyset$.

[0706] In some embodiments, abstract notions of “braking” and “continue forward” behaviors may be assumed. In the structured case, the meanings of these behaviors may be defined based on allowed intervals for lateral and longitudinal accelerations.

[0707] **Definition 43 (Possible Trajectories due to Braking and Normal Driving):** Consider, for example, a vehicle c riding on some road. Given a set of constraints, C , on the behavior of the car, then $T(C, c)$ is the set of possible future trajectories of c if c complies with the constraints given in C . In some embodiments, the set may be a superset. In some embodiments, $T(C_b, c), T(C_f, c)$ may represent the future trajectories due to constraints on braking behavior and constraints on continue forward behavior.

[0708] **Definition 43** may be used to refine the notions of safe distance, dangerous situation, blame time, proper response, and responsibility.

[0709] **Definition 44 (Safe Distance, Dangerous Situation, Blame Time, Proper Response, Responsibility):** The distance between c_0, c_1 driving on an unstructured road is safe if either of the following holds:

1. For some $i \in \{0, 1\}$ then $T(C_b, c_i) \cap T(C_f, c_{1-i}) = \emptyset$ and $T(C_b, c_{1-i}) \cap T(C_f, c_i) \neq \emptyset$
2. $T(C_b, c_0) \cap T(C_b, c_1) = \emptyset$

[0710] In some embodiments, time t is dangerous w.r.t. c_0, c_1 if the distance between them is non safe. The corresponding blame time is the earliest non-dangerous time t_b s.t. during the entire time interval $(t_b, t]$ the situation was dangerous. The proper response of car c_j at a dangerous time t with corresponding blame time t_b is as follows:

1. If both cars were already at a full stop, then c_j can drive away from c_{1-j} (meaning that the absolute value of the angle between its velocity vector and the vector of the difference between c_j and c_{1-j} should be at most θ , where θ is as in **Definition 42**)
2. Else, if t_b was safe due to item (1) above and $j = 1 - i$, then c_j should comply with the constraints of “continue forward” behavior, as in C_f , as long as c_{1-j} is not at a full stop, and after that it should behave as in the case that both cars are at a full stop.
3. Otherwise, the proper response is to brake, namely, to comply with the constraints C_b .

[0711] Finally, in case of a collision, the responsibility is on the vehicle(s) that did not respond properly.

[0712] In some embodiments, the model may account for the braking constraints, C_b , and the continue forward behavior constraints, C_f . As described above, vehicles can predict what other vehicles will do (other vehicles are supposed to drive on their route, and change lateral/longitudinal speed at a bounded rate), and when a vehicle is at a dangerous time, the proper response may be defined with respect to the geometry of the route (“brake” laterally and longitudinally). In some embodiments, the proper response is not defined with respect to the other vehicle from which the host vehicle is at a non-safe distance. For example, conflicts may arise if a vehicle were at a non-safe distance with respect to more than a single other vehicle. Therefore, in some embodiments, when designing the definitions for unstructured scenarios, the two properties will still hold.

[0713] For example, the constraint definitions may rely on a basic kinematic model of vehicles. The vehicle speed may follow the same definition as longitudinal velocity (bounding the range of allowed accelerations). For lateral movements, when a vehicle maintains a constant angle of the steering wheel and a constant speed, it will move (approximately) in a circle. In other words, the heading angle of the car will change at a constant rate, which is called the yaw rate of the car. Let the speed of the car be represented as $v(t)$, the heading angle by $h(t)$, and the yaw rate by $h'(t)$ (as it is the derivative of the heading angle). When $h'(t)$ and $v(t)$ are constants, the car moves on a circle whose “radius” is $v(t)/h'(t)$ (where the sign of the “radius” determines clockwise or counter clockwise and the “radius” is ∞ if the car moves on a line, i.e. $h'(t) = 0$). Thus, the radius may be defined as $r(t) = v(t)/h'(t)$.

[0714] In some embodiments, there may be two constraints on driving. The first is that the inverse of the radius changes at a bounded rate. The second is that $h'(t)$ is bounded as well. The expected braking behavior may be to change $h'(t)$ and $1/r(t)$ in a bounded manner during the response time, and from there on continue to drive on a circle (or at least be at a distance of at most $\epsilon/2$ from the circle). This behavior forms the analogue of accelerating by at most $a_{max,accel}^{lat}$ during the response time and then decelerating until reaching a lateral velocity of zero.

[0715] To make efficient calculations of the safe distance, the superset $T(C_b, c)$ may be described as follows. Without loss of generality, let the blame time be $t = 0$, and assume that at the blame time the heading of c is zero. By the constraint of $|h'(t)| \leq h'_{max}$ then $|h(\rho)| \leq \rho h'_{max}$. In addition, the inverse of the radius at time ρ must satisfy:

$$\frac{1}{r(0)} - \rho r_{max}^{-1'} \leq \frac{1}{r(\rho)} \leq \frac{1}{r(0)} + \rho r_{max}^{-1'}$$

[0716] where $r(0) = v(0)/h'(0)$. Superset $T(C_b, c)$ may be defined as all trajectories such that the initial heading (at time 0) is in the range $[-\rho h'_{max}, -\rho h'_{max}]$, and the trajectory is on a circle whose inverse radius is defined as to the above relationship, and the longitudinal velocity on the circle is as in the structured case. For the continue forward trajectories, the same analysis may apply except that the allowed longitudinal acceleration even after the response time is in $[-a_{max,brake}, a_{max,accel}]$. An illustration of the extreme radiuses is given in Fig. 47H.

[0717] Fig. 47H is an exemplary illustration of the lateral behavior in unstructured scenes. The black line is the current trajectory. The dashed lines are the extreme arcs.

[0718] Finally, these proper responses of the vehicles may satisfy the aforementioned two properties of the proper response for structured scenarios: it is possible to bound the future positions of other vehicles in case an emergency will occur; and the same proper response can be applied even if the vehicle is in a dangerous situation with respect to one or more other vehicles.

[0719] Pedestrians

[0720] The rules for assigning responsibility for collisions involving pedestrians (or other road users) may follow the same ideas described in previous subsections. However, it may not be necessary to adjust the parameters in the definitions of safe distance and proper response, or to specify pedestrians' routes (or possibly unstructured routes) or their priority with respect to vehicles' routes. In some cases, a pedestrian's route is well defined (e.g. a zebra crossing or a sidewalk). In other cases, for example, a residential street, the RSS model may be based on the analysis for unstructured roads. Additionally, unlike vehicles that are restricted to straight lines or arc segments, for pedestrians the change of heading, $|h'(t)|$, may be constrained to assume that in an emergency, after the response time, the pedestrian will continue in a straight line. If the pedestrian is standing, the pedestrian may be assigned all possible lines originating from his current position. The priority may be set according to the type of the road and, possibly, according to the presence and/or locations of traffic lights. For example, on a typical residential street, a pedestrian may have the priority over the vehicles, and it follows that vehicles must yield and be cautious with respect to pedestrians. In contrast, there may be roads with a sidewalk where the vehicles may not need to assume that a pedestrian on the sidewalk may suddenly run into the road. In such situations, for example, cars have the priority. Another example is a zebra crossing with a traffic light, where the priority is set dynamically according to the light. In some embodiments, and as previously described, priority is given not taken, hence even if pedestrians do not have priority, if the pedestrian entered the road at a safe distance, cars must brake and let them pass.

[0721] For example, take a scenario in which a pedestrian that stands on a residential road. The pedestrian is assigned to all routes obtained by rays originating from its current position. Her safe longitudinal distance with respect to each of these virtual routes may be short. For example, setting a delay of 500 ms, and maximal acceleration and braking of 2 m/s^2 , yields that the pedestrian's part of the safe longitudinal distance is 50cm. It follows that a vehicle may be in a kinematic state such that, if it

applies a proper response (acceleration for ρ seconds and then braking) it will remain outside of a radius 50cm around the pedestrian.

[0722] In a second example, a pedestrian may be standing on the sidewalk in front of a zebra crossing. The pedestrian may have a red light, and a vehicle approaching the zebra crossing may have a green light. In this example, the vehicle route has the priority, hence the vehicle can assume that the pedestrian will stay on the sidewalk. If the pedestrian enters the road while the vehicle is at a safe longitudinal distance (with respect to the vehicle's route), then the vehicle must brake. However, if the pedestrian enters the road while the vehicle is not at a safe longitudinal distance, and as a result the vehicle hits the pedestrian, then the vehicle is not responsible. It follows that in this situation, the vehicle can drive at a normal speed, without worrying about the pedestrian.

[0723] A third example is when a pedestrian runs on a residential road at 10 km per hour (which is ≈ 2.7 m/s). The possible future trajectories of the pedestrian form an isosceles triangle. Using the same parameters as in the first example, the height of this triangle is roughly 15m. It follows that cars should not enter the pedestrian's route at a distance smaller than 15m. But, if the car entered the pedestrian's route at a distance larger than 15m, and the pedestrian did not stop and, as a result, crashed into the car, then the responsibility is on the pedestrian.

[0724] Cautiousness with respect to occlusion

[0725] Human sensing capabilities are limited, sometimes because of, for example, a subconscious decision to focus on a specific part of the road, sometimes because of carelessness, and sometimes because of physical limitations (e.g., it is impossible to see a child hidden behind a parked car). While advanced automatic sensing systems may also suffer from limited sensing due to physical occlusions or range of sensor detection. For example, a child may be out of sight behind a parked car, a building or fence may block sight of traffic approaching an intersection, or blind spots may be created by other vehicles on a highway.

[0726] Fig. 47I is an exemplary scenario in which c_0 is trying to exit a parking lot and merge into a route but cannot see whether there are cars approaching the merge point from the left side of the street. If this is an urban, narrow street, with a speed limit of 30 km/h, a human driver's behavior may be to slowly merge onto the road, obtaining more and more field of view, until sensing limitations are eliminated. Without making additional assumptions, this behavior may lead to a collision. Indeed, if c_1 is driving very fast, then at the moment that it enters the view range of c_0 , c_1 might reveal that the situation is already dangerous for more than ρ seconds and cannot respond properly. **Definition 45** provides a definition for the first time an occluded object is exposed to the sensor; after its exposure, the system can handle the situation as it would any other object the sensor can detect.

[0727] **Definition 45 (Exposure Time):** The Exposure Time of an object may be the first time at which it is detected by the sensor.

[0728] Thus, using **Definition 45**, it may be possible to set an upper bound on the reasonable speed of road users. In the previously described example, it is reasonable to assume that c_1 will not drive faster than 60 km/h (twice the allowed speed). With this assumption, c_0 can approach the merge point

slowly enough so as to make sure that it will not enter a dangerous situation with respect to any car whose speed is at most 60 km/h. This leads to the following definition.

[0729] **Definition 46 (Responsibility due to Unreasonable Speed):** Consider, for example, an accident between c_0 , c_1 , where, before the blame time, there was a safe lateral distance, and where c_1 was driving unreasonably fast, namely, its average velocity from the exposure time until the blame time was larger than v_{limit} , where v_{limit} is a parameter associated with the position of c_1 on the map and on one or more road conditions. In this exemplary situation, the responsibility is solely on c_1 .

[0730] In some embodiments, application of these definitions allows c_0 to exit the parking lot safely, in the same manner a human driver does. Based on the above definition, c_0 may check the situation where there is a vehicle at the edge of the occluded area whose velocity is v_{limit} . Thus, this encourages c_0 to drive slower and further from the occlusion, thus slowly increasing its field of view and allowing for safe merging into the street.

[0731] This responsibility definition may apply to a variety of cases as it may specify what can be occluded (a potentially fast car cannot be occluded between two closely parked cars, but a child or pedestrian can). This responsibility definition may also specify what is unreasonably fast for an occluded object or person (a child's v_{limit} is much smaller than that of a car).

[0732] Occluded pedestrians in a residential area

[0733] In one exemplary situation, a child may run out from behind a parked car. If the road is in a residential area, the child's route has priority, and hence a vehicle should adjust its speed such that, if the child emerges from behind the parked car then there will be no accident. By limiting the maximal speed of the child, the vehicle may adjust speed according to the worst-case situation in which the child will be exactly at the maximal speed.

[0734] However, in some embodiments, the above approach may yield a too defensive behavior. Indeed, consider a typical scenario when a vehicle is driving next to a sequence of parking cars. A child, that runs into the road (e.g. chasing a ball) at a speed of 10 km/h, might be occluded by a parked car, where the vehicle's camera will sense the child only when the longitudinal distance between the vehicle and the pedestrian is smaller than 0.3 m. Letting the longitudinal safe distance of the child's route be around 15 m, which is typically more than the width of the residential road. Therefore, the vehicle is already at a non-safe longitudinal distance (with respect to the child's route) and must be able to stop before entering the child's route. However, the braking distance of a car driving at 1 m/s (with reasonable setting of parameters) may be more than 0.4 m. It follows that the vehicle should drive slower than 1 m/s in this scenario, even if driving at a lateral distance of, for example, 5 m, from the parking cars. This may be too defensive and does not reflect normal human behavior.

[0735] In some embodiments, assumptions may be made such that the vehicle does not overcompensate and drive too defensively. For example, a human driver may rely on the law. While autonomous vehicles will follow the speed limit, some embodiments may require more robust notions of safety. Below is an exemplary responsibility definition for this scenario, that leads to a safer driving while enabling reasonably fast driving.

[0736] **Definition 47 (Responsibility for Collision with Occluded Pedestrian):** Consider, for example, a collision between a pedestrian and a vehicle at a residential road. The vehicle is not responsible for the accident if the following two conditions hold:

- Let t_e be the exposure time. The vehicle did not accelerate at the time interval $[t_e, t_e + \rho)$, and performed a longitudinal brake of at least $a_{\min, \text{brake}}$ from $t_e + \rho$ until the accident or until arriving to a full stop.
- The averaged velocity of the vehicle from the exposure time until the collision time was smaller than the averaged velocity of the pedestrian at the same time interval (where averaged velocity is the total distance an agent moved through a time interval divided by the length of the time interval).

[0737] The first condition represents that the velocity of the vehicle may be monotonically non-increasing from the exposure time until the collision time, and the second condition tells us that the averaged velocity of the vehicle is smaller than that of the pedestrian. As implied by the proof of the lemma below, under some reasonable assumptions on the pedestrian maximal acceleration, the responsibility definition tells us that at the collision time, it may be expected that either the velocity of the car will be significantly smaller than that of the pedestrian, or that both the vehicle and the pedestrian will be at a very slow speed. In the latter case, the damage should not be severe and in the former case, it complies with “common sense” that the pedestrian is to be blamed for the accident.

[0738] The lemma below shows the allowed driving speed that guarantees to never be responsible for a collision with an occluded pedestrian.

[0739] **Lemma 14:** *Suppose that a vehicle located at a lateral distance of x and longitudinal distance of y from an occluded spot. Then, if its lateral speed is zero and its longitudinal speed is smaller than*

$$v_{\max}(x, y) = a_{\min, \text{brake}} \left(-\rho + \sqrt{\rho^2 + \frac{x^2 + y^2}{y a_{\min, \text{brake}}}} \right)$$

[0740] *it can guarantee to not make an accident of the vehicle's responsibility. In particular, a constant driving speed of*

$$v_{\max}(x) = \min_{y>0} v_{\max}(x, y) = a_{\min, \text{brake}} \left(-\rho + \sqrt{\rho^2 + \frac{2x}{a_{\min, \text{brake}}}} \right)$$

[0741] *is possible.*

[0742] **Proof:** Suppose that at the exposure time, the pedestrian is located at position $(0, 0)$ and the car is located at position $(x, -y)$. Suppose the collision is at the point (x, y_c) and let (x, y_s) be the point in which the car will come to a full stop (if there will be no collision). Define by t_c the collision time and by t_s the time the vehicle will reach a full stop. Note that if there is a collision of the vehicle's blame at point (x, y_c) , then:

$$x^2 + y_c^2 \leq (y_c + y)^2 \Rightarrow x^2 \leq y^2 + 2yy_c \Rightarrow \frac{x^2 - y^2}{2y} \leq y_c$$

[0743] But, since $y_s \geq y_c$, this also implies that

$$y_s \geq \frac{x^2 + y^2}{2y} \Rightarrow x^2 + y_s^2 \leq (y_s + y)^2$$

[0744] which implies that there can be a collision of the vehicle's blame at the braking time.

So, to make sure that there is no collision of the vehicle's blame, verify that there is no such collision at

5 the braking time, that is, that the following condition holds: $y_s \frac{x^2 - y^2}{2y}$.

[0745] Since $y_s = -y + v\rho + 0.5 v^2/a_{min,brake}$, rearrangement yields:

$$\frac{v^2}{2a_{min,brake}} + v\rho - \frac{x^2 + y^2}{2y} < 0$$

[0746] This condition holds for the interval $v \in [0, v_{max}(x, y)]$ where $v_{max}(x, y)$ is the right root

10 of the quadratic equation, namely:

$$v_{max}(x, y) = a_{min,brake} \left(-\rho + \sqrt{\rho^2 + \frac{x^2 + y^2}{ya_{min,brake}}} \right)$$

[0747] Since $f(y) = (x^2 + y^2)/y$ is minimized at $y = x$ yielding that the maximal constant speed of the vehicle is:

$$v_{max}(x) = a_{min,brake} \left(-\rho + \sqrt{\rho^2 + \frac{2x}{a_{min,brake}}} \right)$$

15 [0748] RSS For Advanced Driving Assist Systems

[0749] Active safety systems may help make the road safer. For example, an Autonomous Emergency Braking (AEB) system may slow down, or may stop the car, when it detects that the car is likely to hit the car in front of it. One problem arising with AEB systems is that they are tuned to have close to zero false positives (a false positive is when the system detects a car at a dangerous position, while in reality no car is at a dangerous position). The price for extremely low false positive rate is a much higher false negative rate (a false negative is when the system fails to detect a car which is at a dangerous position, and hence the system does not respond to the dangerous situation). Thus, there exists a need for constructing a better AEB system relying on a generalization of the Responsible Sensitive Safety (RSS) concept. In some embodiments, RSS accounts for the tradeoff between false positive / false negative better than existing systems. The RSS may use a braking profile that can tolerate a higher degree of false positives, and therefore is less dangerous.

[0750] Recall the RSS Safe Distance (Definition 27). RSS's proper response does not depend on the underlying driving policy. This definition may be embedded on top of any driving policy. For example, RSS may be embedded on top of a human driving policy. In this case, whenever a human driver brings the car to a non-safe distance, the RSS-ADAS system may apply braking so as to bring the car back to a safe distance.

[0751] However, applying a strong brake without any warning might be dangerous. Consequently, existing AEB systems are tuned so as to have an extremely low probability of false

positives. The price for extremely low false positive rate is a much higher false negative rate (meaning, the system fails to detect a car which is at a non-safe distance). The generalization of RSS, discussed above, may be used to specify a particular member of this RSS family that mitigates the danger of false positives in an alternative way.

5 [0752] Generalized RSS

[0753] In the original definition of RSS, the rear car is assumed to accelerate during the response time and then to brake until reaching a full stop. This is an example of a *braking profile*.

[0754] More generally:

10 [0755] **Definition 48 (Braking profile):** A braking profile, B , may be a mapping from initial kinematic state of the car (mainly, initial velocity, v_0 , and acceleration, a_0) to a pair T_b, v , s.t. $v : [0, \infty) \mathbb{R}$ is the future velocity of the car and $T_b < 0$ is the first time in which $v(t) = 0$.

[0756] For example, the braking profile applied by the front car in the original definition of RSS may be defined by $T_b = \frac{v_0}{a_{\max, \text{brake}}}$ and $v(t) = \max \{ v_0 - ta_{\max, \text{brake}}, 0 \}$, such that:

$$v(t) = \begin{cases} v_0 + ta_{\max, \text{accel}} & \text{if } t \leq \rho \\ v_0 + \rho a_{\max, \text{accel}} - (t - \rho)a_{\min, \text{brake}} & \text{if } t \in (\rho, T_b) \\ 0 & \text{otherwise} \end{cases}$$

15 [0757] **Definition 49 (Generalized RSS Safe Distance with respect to Braking Profiles: $B_f ; B_r$):** A longitudinal distance between a car c_r that drives behind another car c_f , where both cars are driving at the same direction, is safe with respect to braking profiles $B_f ; B_r$ if in case the front car applies braking profile B_f and the rear car applies braking profile B_r , then the cars will reach a full stop without colliding.

[0758] The proper response may be defined as follows:

20 [0759] **Definition 50 (Proper response):** Suppose the first time that the distance between c_f and c_r is non-safe is t_0 , and without loss of generality set $t_0 = 0$. Then, the proper response for the front car is to have its velocity at least $v_f(t)$, where v_f is the velocity defined by B_f , and the proper response for the rear car is to have its velocity at most $v_r(t)$, where v_r is the velocity defined by B_r .

[0760] If both cars apply proper response then there will be no accident.

25 [0761] Jerk-Bounded Braking Profile

[0762] The jerk-and-acceleration-bounded braking profile may be denoted B_j . The vehicle may start decreasing acceleration linearly (with slope j_{\max}), until reaching a max brake parameter (denoted $a_{\min, \text{brake}}$), and then continue to brake with a constant deceleration until reaching zero velocity. The following closed form formulas for may be used for calculating this braking profile.

30 [0763] **Braking with a Constant Jerk:** Suppose the vehicle starts braking at jerk of j_{\max} . Suppose also that the vehicle can come to acceleration 0 immediately, so below, assume $a_0 \leq 0$. Then, the dynamics of the car is as follows:

$$a(t) = a_0 - j_{\max} t$$

$$v(t) = v_0 + \int_{\tau=0}^t a(\tau) d\tau = v_0 + \left[a_0 \tau - \frac{1}{2} j_{\max} \tau^2 \right]_0^t = v_0 + a_0 t - \frac{1}{2} j_{\max} t^2$$

$$x(t) = x_o + \int_{\tau=0}^t v(\tau) d\tau = x_o + \left[v_o \tau + \frac{1}{2} a_o \tau^2 - \frac{1}{6} j_{max} \tau^3 \right]_0^t = x_o + v_o t + \frac{1}{2} a_o \tau^2 - \frac{1}{6} j_{max} t^3$$

[0764] **Braking profile:** Based on these equations, braking distance will be defined as follows. Let T be the first time in which either $a(T) = a_{min,brake}$ or $v(T) = 0$, that is, $t =$

5 $\min\{T_1, T_2\}$ where $T_1 = \frac{a_o + a_{min,brake}}{j_{max}}$, $T_2 = \frac{a_o + \sqrt{a_o^2 + 2j_{max}v_o}}{j_{max}}$

[0765] The time for reaching full brake is

$$T_b = \begin{cases} T_2 & \text{if } T = T^2 \\ \frac{v_o + a_o T - \frac{1}{2} j_{max} T^2}{a_{min,brake}} & \text{otherwise} \end{cases}$$

[0766] And the speed is as follows

$$v(t) = \begin{cases} v_o + a_o T - \frac{1}{2} j_{max} T^2 & \text{if } t \leq T \\ v_o + a_o T - \frac{1}{2} j_{max} T^2 - (t - T) a_{min,brake} & \text{if } t \in (T, T_b) \\ 0 & \text{otherwise} \end{cases}$$

10 [0767] **Braking distance:** Finally, the car will move the following distance until reaching a full stop:

$$\left[v_o T + \frac{1}{2} a_o T^2 - \frac{1}{6} j_{max} T^3 \right] + \frac{\left(v_o + a_o T - \frac{1}{2} j_{max} T^2 \right)^2}{2 a_{min,brake}}$$

[0768] Safe distance with respect to B_f and B_j

[0769] Suppose that the front car brakes with braking profile B_f (braking with a constant deceleration of $a_{max,brake}$), and the rear car relies on braking profile B_j . Then, the safe distance formula is

$$\left[\left[v_o T + \frac{1}{2} a_o T^2 - \frac{1}{6} j_{max} T^3 \right] + \frac{\left(v_o + a_o T - \frac{1}{2} j_{max} T^2 \right)^2}{2 |a_{min,brake}|} - \frac{v_f^2}{2 |a_{min,brake}|} \right]_+$$

[0770] The ADAS RSS System

[0771] As soon as the distance becomes non-safe (according to the jerk-based formula), the system may alert the driver and may start to brake with a jerk of j_{max} until one of the following happens: (1) the car stops (2) the distance becomes non-safe (3) the driver turns off the RSS system (by a button or some other mechanism).

[0772] The same proof of RSS may be adjusted to show that if the driver did not turn-off the system, then there can be no accident of the ego vehicle's responsibility. As an example, if all vehicles follow the RSS rules then there will be no accidents at all (due to driving policy errors, sensing mistakes or malfunctioning can still happen). Therefore, if a regulatory body adopt the ADAS RSS system it may reduce the car accidents significantly.

[0773] A better handling of false positives

[0774] The same hardware that implements AEB may be adjusted to implement ADAS RSS. Because the braking profile is jerk bounded, the danger of false positives is not very severe (the driver is not surprised, and someone behind the rear car will also not be surprised because of the bound on the jerk). Therefore, it is possible to tune the system to have less false negatives.

5 [0775] Navigation with a Safe Longitudinal Distance

[0776] As described above, RSS may be effective for determining a safe distance between a host vehicle and another object (e.g., one or more target vehicles or VRUs). RSS provides a set of mathematical formulae that can be used to ensure that a vehicle implementing RSS will not be involved in an accident caused by it. As discussed above, RSS may be applied to maintain a safe distance between
10 two vehicles traveling in the same direction. In other embodiments, however, RSS may also be used to maintain a safe distance between two vehicles traveling in opposite directions. For example, two vehicles may be driving towards each other during a lane change or while navigating a parking lot. In some instances, a vehicle may be traveling on the wrong side of the road. RSS may be used to ensure that a host vehicle maintains a safe distance accounting for both vehicles and braking (and/or steering) if needed
15 such that the vehicles come to a complete stop (or otherwise avoid one another) without causing a collision.

[0777] Fig. 48A illustrates an example scenario with two vehicles traveling in opposite directions consistent with the disclosed embodiments. In this scenario, a host vehicle 4802 may be traveling along a lane of a two-lane road, as shown in Fig. 48A. A target vehicle 4804 may be traveling in
20 the opposite direction of host vehicle 4802 in a lane occupied by host vehicle 4802. For example, target vehicle 4804 may be passing a second oncoming vehicle 4806 and may enter the lane of travel of host vehicle 4802 in order to complete the passing maneuver. Thus, host vehicle 4802 and target vehicle 4804 may be traveling toward each other in opposite directions. Various other scenarios may exist where two vehicles are traveling in opposite directions. For example, host vehicle 4802 may be performing the
25 passing maneuver and may enter the lane of travel of target vehicle 4804. The two vehicles may be traveling towards each other in other situations, such as in a parking lot, when they are both proceeding through an intersection, or any other similar scenario. The present disclosure is not limited to any particular scenario in which a direction of travel of a target vehicle is toward the host vehicle.

[0778] Host vehicle 4802 may be configured to perform a planned navigational action if it
30 allows host vehicle 4802 to maintain a safe longitudinal distance from vehicle 4804, such as an RSS safe distance, as discussed above. In the scenario shown in Fig. 48A, for example, the planned navigational action may be to maintain a forward heading direction or to maintain a throttle action of host vehicle 4802. In other embodiments, the planned navigation action may be one that would cause host vehicle 4802 and target vehicle 4804 to travel towards each other in opposite directions. For example, the
35 navigational action may include a passing maneuver, in which host vehicle 4802 attempts to pass another vehicle and enters a lane occupied by target vehicle 4804. Host vehicle 4802 may be configured to perform the passing maneuver if a safe distance would be maintained between host vehicle 4802 and target vehicle 4804. The navigation action may include other maneuvers as well, such as a lane change

maneuver, a merge maneuver, a parking lot navigation maneuver (e.g., entering a parking spot, leaving a parking spot, parallel parking, maneuvering between parking rows, etc.), an intersection maneuver (e.g., turning across oncoming traffic, proceeding through a yellow light, etc.), or the like.

[0779] In the example scenario of Fig. 48A, the direction of target vehicle 4804 is toward host vehicle 4802 in a fully “head-on” situation. The present disclosure may equally apply to situations in which the direction of target vehicle 4804 is toward host vehicle 4802 at an angle. Accordingly, the term “toward” as used herein should include any situation in which a substantial component of the target vehicle’s traveling direction is in the direction of the host vehicle. The term “toward” may also refer to a situation in which the frontal corridor of two or more vehicles at least partially overlap.

[0780] Fig. 48B illustrates an example scenario of a target vehicle 4804 traveling towards a host vehicle 4802 consistent with the disclosed embodiments. Host vehicle 4802 may be traveling in a direction 4812 and target vehicle 4804 may be traveling in a direction 4814 that intersects direction 4812. Directions 4812 and 4814 may intersect at an angle, Θ , as shown. In a “head on” situation, where host vehicle 4802 and 4804 are traveling directly at each other, angle Θ may be equal or substantially equal to zero. Target vehicle 4804 may be considered to travel toward host vehicle 4804 at other angles as well. For example, in some embodiments, direction 4814 of target vehicle 4804 may be considered to be toward host vehicle 4802 as long as angle Θ is less than 90 degrees. In other embodiments, angle Θ may be required to be less than another angle (e.g., 60 degrees, 45 degrees, 30 degrees, etc.). Direction 4814 may also be defined based on longitudinal and lateral components (x and y , respectively) in relation to host vehicle 4802. Direction 4814 of target vehicle 4804 may be considered to be toward host vehicle 4802 as long as x is some nonzero component of direction 4814 in the direction of host vehicle 4802. In some embodiments, a relationship between component x and component y may also be defined. For example, direction 4814 of target vehicle 4804 may be considered to be toward host vehicle 4802 as long as component x is greater than component y (representing an angle Θ of 45 degrees). Accordingly, it is to be understood that the longitudinal safety distance is not limited to direct “head on” situations and may be at various other angles.

[0781] As discussed above, the RSS safe distance may correspond to a distance between a host vehicle and a target vehicle at which the host vehicle may come to a stop without colliding with the target, assuming certain conditions. The RSS safe distance may include a calculation of the minimum distance that it would take the vehicles to stop if braking at predetermined braking rates and after an assumed behavior during corresponding reaction times for the vehicles. For example, the RSS safe distance accounts for the possibility that the two cars travel toward one another both with a maximum acceleration during a reaction time before each braking at a predetermined braking level to come to a complete stop. Taking into account the vehicles’ maximum acceleration during reaction times allows navigation according to the RSS distance to ensure no collisions occur even in worst case scenarios.

[0782] Fig. 49 illustrates an example of host vehicle maintaining a safe longitudinal distance consistent with the disclosed embodiments. Target vehicle 4904 may be traveling in a direction toward host vehicle 4902. A processing device of host vehicle 4902 (e.g., processing unit 110) may determine a

planned navigational action for accomplishing a navigational goal of the host vehicle. Host vehicle 4902 may be configured to perform the planned navigational action if an RSS safe longitudinal distance can be maintained between host vehicle 4902 and target vehicle 4904. As discussed above, the planned navigational action may include various maneuvers, such as a lane change maneuver, a merge maneuver, a passing maneuver, a maintain forward heading maneuver, a parking lot navigation maneuver, or a maintain throttle action, or other similar navigational actions. The processing device of host vehicle 4902 may determine a next-state distance 4910 between host vehicle 4902 and target vehicle 4904 corresponding to a distance between the vehicles that would result if the planned navigational action was taken. The processing device may also determine a stopping distance 4912 in which vehicle 4902 can be stopped and a stopping distance 4914 in which vehicle 4904 can be stopped. The processing device of host vehicle 4902 may implement the planned navigational action if the determined next-state distance 4910 is greater than a sum of stopping distance 4912 and stopping distance 4914, thereby indicating the vehicles can both come to a complete stop to avoid a collision.

[0783] The stopping distances 4912 and 4914 may be calculated in any manner described above in relation to RSS. For example, the distance needed to come to a stop may be calculated using a vehicle braking rate, a vehicle maximum acceleration capability, and a current speed of the vehicle. In some embodiments, the current speed of the host vehicle may be determined based on an output of one or more sensors. For example, the current speed may be determined from an output of an accelerator, a LIDAR system, a RADAR system, a GPS unit, or the like. As another example, the current speed of the host vehicle may be determined by analyzing one or more images (e.g., based on a scaling rate change of a fixed object detected in two or more images).

[0784] In some embodiments, the disclosed embodiments may assume that a vehicle will brake at its maximal braking capability to avoid a collision. For example, the braking rate may be based on based on at least one recognized characteristic of the vehicle. For host vehicle 4902, the maximal braking rate may be known or estimated by the system based on characteristics of the car, previous braking maneuvers, or the like. For a target vehicle, the maximal braking capability may be determined based on at least one characteristic of the target vehicle. The at least one characteristic may be any characteristic that is known to correlate with a braking capability or that may correlate with a braking capability. In some embodiments, the recognized characteristic of the target vehicle may include a vehicle type, which may be determined based on analysis of images containing representations of the target vehicle. The vehicle type may be, for example, a general category to which the target vehicle belongs (e.g., full-size sedan, compact car, SUV, cross-over SUV, motorcycle, etc.) or a more particular category or sub-category associated with the target vehicle. In some embodiments, the recognized characteristic of the target vehicle may include a vehicle model, brand name, or other classifier of the target vehicle (e.g., Toyota Prius®, BMW X5®, etc.). The processing device may be programmed to assume that the target vehicle has a maximum braking capability corresponding with that of a known braking capability of that vehicle type or a class of vehicles. The processing device may use the vehicle type to, for example, access a database containing known braking capabilities of each vehicle type. For example, if the detected

vehicle type is BMW X5®, the processing device may look up the reported braking capabilities of a BMW X5® from a database. If there are multiple maximum braking capabilities reported for a vehicle model (for example, a BMW X5® may have different braking capabilities depending on whether it has the xDrive40i or xDrive50i trim level or depending on the year of its manufacture), the processing device
5 may be programmed to assume the target vehicle has the best maximum braking capability reported for that vehicle model. The recognized characteristic of the target vehicle may also include an indicator of whether the target vehicle is an autonomous vehicle, partially autonomous vehicle, or whether the target vehicle is human-operated. Such information may be derived from analysis of captured images (e.g., images acquired from one or more cameras onboard the host vehicle) or may be broadcast from the target
10 vehicle via wireless communication.

[0785] In some embodiments, the recognized characteristic of the target vehicle may include a vehicle size. The vehicle size may be a relative size, a predetermined size, or the like. The vehicle size may correspond to a physical size of the vehicle and/or an estimated weight of the vehicle. For example, the processing device may determine that the target vehicle is larger or smaller than the host vehicle. In
15 another example, the processing device may be programmed to classify the target vehicle into one or a plurality of predetermined vehicle sizes, each of which may include a range of sizes (e.g., category 1 = less than 1,000 kg; category 2 = 1,000–1500 kg; category 3 = 1,500-1,800 kg; and so forth). Some vehicles may contain an indication of its size, which may be used to determine a vehicle size of the target vehicle. For example, the target vehicle may be a GMC Sierra® and may include an indication or model
20 name on the exterior of the vehicle associated with its size, such as 1500, 2500, 2500HD, and so forth.

[0786] Applying maximal braking, especially when such application is sudden, may be regarded as an extreme response that should be reserved for cases when it is not avoidable. Under certain circumstances, maximal braking can be uncomfortable to passengers, can implicate trailing vehicles, can excessively wear down vehicle hardware (tires, brake pads etc.), etc. Accordingly, the stopping distance
25 for the vehicle may assume a predetermined submaximal braking rate that is less than a maximum braking rate capability of the host vehicle. For example, the submaximal braking rate may assume a gradual application of the brake from an initial braking force (that is less than the maximum braking capability of the vehicle) to the maximum braking capability of the vehicle rather than an immediate application of the maximum braking capability. In some embodiments, the submaximal braking rate may
30 be an ideal or statistically acceptable braking rate. For example, data relating to the braking rate at which most humans come to a stop may be used to determine submaximal braking rate. In some embodiments, the submaximal braking rate may be set according to a user preference. For example, a first passenger may provide a first the submaximal braking rate that is comfortable to him or her, a second passenger may provide a second the submaximal braking rate that is comfortable to him or her, and so forth. In
35 some embodiments, a processing device may determine the submaximal braking rate based on, for example, one or more sensed road conditions (e.g., dirt road, road traction properties, etc.), traffic conditions (congestion level, average distance between vehicle around the host vehicle, rate of cut-ins in front of the host vehicle, etc.), one or more vehicle characteristics, one or more sensed weather

conditions, or the like. For example, a processing device may determine a submaximal braking rate that provides the maximum safety and comfort on a snowy road, a submaximal braking rate that provides the maximum safety and comfort on a dry highway, and so forth. It will be appreciated, that while implementing the submaximal braking rate can help increase comfort and have other positive impacts on the behavior of the host vehicle, it can also have less desirable effects, in particular, if the submaximal braking rate is small, so that braking may become "relaxed" or not as strong. For example, using a small submaximal braking rate may have a strong negative effect on efficiency, by causing the vehicle to always keep an exaggerated distance from the vehicle ahead. This may cause an annoyance to other road users and prevent the host vehicle from effectively negotiating complex scenarios such as merging and driving in congested areas. Thus, in some examples, a submaximal braking rate that balances performance or efficiency with comfort (while maintaining RSS safety) may be set, selected or determined.

[0787] Consistent with this disclosure, the submaximal braking rate may be a constant. For example, the submaximal braking rate may be a braking rate that corresponds with a deceleration of 2 m/s^2 , 10 m/s^2 , or any other constant deceleration. In some embodiments, the submaximal braking rate may be a constant that is proportional to the maximum braking rate capability of a host vehicle. For example, the submaximal braking rate may be 20%, 33%, 50%, 70%, or any other percentage of the maximum braking capability of the host vehicle. In some embodiments, the submaximal braking rate may be a linear function of the current speed of the host vehicle, such that the brake is gradually applied from the minimum braking capability of the host vehicle up to, at most, the maximum braking capability of the host vehicle. For example, a submaximal braking rate may be determined that allows the host vehicle to come to a stop from its current speed without reaching the maximum braking capability of the vehicle. As another example, a submaximal braking rate may be determined that eventually reaches the maximum braking capability of the host vehicle such that the host vehicle brakes at less than maximum braking capability for a first portion of a stopping maneuver and brakes at its maximum braking capability for a second portion the stopping maneuver. In some embodiments, the submaximal braking rate may be a different function of the current speed of the vehicle. For example, the submaximal braking rate may be a braking rate that is exponentially or logarithmically applied until the vehicle comes to a stop.

[0788] In some embodiments the submaximal braking rate may be determined based on a blame or responsibility determination associated with host vehicle 4902 or target vehicle 4904. From a responsibility perspective, host vehicle 4902 may have an obligation to brake, but at a lower rate than target vehicle 4904. For example, target vehicle 4904 may be traveling in an incorrect direction for a given lane or roadway, as shown in Fig. 48A, and thus may have a greater obligation to brake. Accordingly, the host vehicle braking rate may be less than the target vehicle braking rate, or vice versa.

[0789] As discussed above, an RSS distance may include be dependent upon the maximum possible acceleration capability that the vehicle may experience during a period between a sensed event and a vehicle's reaction to a sensed event (i.e., a reaction time associated with the host vehicle). Accordingly, the stopping distance may also include a vehicle acceleration distance that corresponds to a distance the vehicle can travel during a vehicle reaction time at the vehicle's maximum acceleration

capability, starting from the determined current speed of the vehicle. For example, the disclosed embodiments may account for a delay between the period in which host vehicle 4902 or target vehicle 4904 begins applying the brake or making another navigational maneuver. For example, there may be several milliseconds (or more or less) between the time host vehicle 4902 detects that target vehicle 4904 is traveling towards host vehicle 4902. The processing device may account for the acceleration that may occur during response period ρ , corresponding to the vehicle's reaction time. Response period ρ may correspond to a known characteristic of a vehicle, similar to the braking rate, or may be an assumed value for all cars (e.g., a certain reaction time for human-operated vehicles and another value (possibly different) for autonomous vehicles). For example, let v_1 be the current velocity of the vehicle. The stopping distance for a vehicle may include a first term corresponding to a distance the vehicle travels as a result of a maximum acceleration during the reaction time. The stopping distance may also include a distance the vehicle requires to reduce its velocity to 0, at a predetermined braking rate (e.g., a maximal braking rate or a submaximal braking rate), from its final velocity at the end of the reaction time period.

[0790] As noted, the velocity of the vehicle used in the RSS and/or stopping distance determination may correspond to a maximum velocity of the host vehicle after accelerating for its response period, ρ (e.g., at its maximum acceleration capability). For example, let v_1 be a current velocity of the vehicle, let ρ be the response period, and let $a_{max,accel}$ be the maximum acceleration of the vehicle, then the maximum velocity $v_{max,\rho}$ of the vehicle after the response period ρ is:

$$v_{max,\rho} = v_1 + \rho a_{max,accel}$$

[0791] Accordingly, $v_{max,\rho}$ may be used in determining the stopping distance of host vehicle 4902 and/or target vehicle 4904. For example, the processing device may assume the maximal or submaximal braking rate is applied starting at a speed of $v_{max,\rho}$ in determining the stopping distance of the vehicle.

[0792] In some embodiments, the processing device of host vehicle 4902 may be configured to implement the planned navigational action if the determined next-state distance is greater, by at least a predetermined minimum distance (e.g., 0 or any suitable distance greater than 0, such as 0.5 m, 1 m, 2 m, etc.), than the sum of the stopping distance 4912 for host vehicle 4902 and the stopping distance 4914 of target vehicle 4904. If the RSS safe longitudinal distance includes a non-zero minimum approach distance component, then in the safe distance determination, host vehicle 4902 would come to a stop at least the minimum distance from the target vehicle. The minimum distance may be predetermined (e.g., a minimum distance of 0.5 m, 1 m, 2m, etc.). The minimum approach distance may be determined based on any set of factors consistent with this disclosure. In some cases, the minimum approach distance may be based on vehicle velocity. For example, a minimum approach distance may be set higher for higher rates of sensed velocity relative to either host vehicle 4902 or target vehicle 4904. The minimum approach distance may also depend on a sensed weather condition (e.g., where snow or water is sensed on a road, or where a sensed traction indicator suggests less than optimal traction, the minimum approach distance may be increased to provide an additional safety buffer. The minimum approach distance may also be set based on a user's preferences, etc..

[0793] Figs. 50A and 50B provide flowcharts depicting an exemplary process 5000 for maintaining a safe longitudinal distance consistent with the disclosed embodiments. Process 5000 may be performed by at least one processing device, such as processing unit 110, described above. Process 5000 is exemplary, and one of ordinary skill in the art having the benefit of this disclosure may understand that process 5000 may include additional steps, exclude certain steps, or may be otherwise modified in a manner consistent with this disclosure.

[0794] Process 5000 may include a step 5002 for receiving at least one image representative of an environment of the host vehicle. The at least one image may be received from an image capture device, which may be any image capture device consistent with this disclosure, including image capture device 122. In some embodiments, the at least one image may be an image obtained from any of a camera, a RADAR, a LIDAR, or any other device from which an image may be obtained, whether optical or otherwise. There may some period of delay between the time when the image is captured and time when the processing device receives the image. Further, there may a period of delay between the time when an event occurs and the imaging device captures an image of the event. For example, if a target vehicle enters a lane in front of a host vehicle, there may be a short period (e.g., a millisecond, 2 milliseconds, 5 milliseconds, etc.) between when the target vehicle maneuvers into the lane and when the imaging device captures (an processes) an image of the target vehicle.

[0795] Process 5000 may include a step 5004 for determining a planned navigational action for the host vehicle. Consistent with this disclosure, the at least one processing device may be programmed to determine a planned navigation action for accomplishing a navigational goal of the host vehicle. The navigation action may be determined based on at least one driving policy. The planned navigation action and/or the at least one driving policy may be any consistent with this disclosure, including those discussed above. For example, the planned navigation action may include at least one of a lane change maneuver, a merge maneuver, a passing maneuver, a maintain a forward heading maneuver, a parking lot navigation maneuver, or a maintain throttle action.

[0796] Process 5000 may include a step 5006 for analyzing the at least one image to identify a target vehicle in the environment of the host vehicle. The at least one image may be an image received from an image capture device, such as image capture device 122. The at least one image may be the one or more images received as part of step 5002 of process 5000.

[0797] Process 5000 may include a step 5008 for determining a next-state distance associated with the planned navigation action. Consistent with this disclosure, the processing device may be programmed to determine a next-state distance between the host vehicle and the target vehicle that would result if the planned navigation action was taken. The next-state distance may be calculated by any means disclosed herein, including the RSS safe distance described above. For example, if the planned navigation action is an acceleration of the host vehicle, the next-state distance may be an expected distance between the host vehicle and leading target vehicle in front of the host vehicle after the navigational action is taken.

[0798] Process 5000 may include a step 5010 for determining a host vehicle braking rate, a host vehicle maximum acceleration capability, and a current speed of the host vehicle. Each of the factors, (i.e., the vehicle braking rate, the maximum acceleration capability, and the speed of the host vehicle) may be determined by any means consistent with this disclosure. The vehicle braking rate may represent a maximal braking rate associated with the vehicle, or may be a submaximal braking rate, as described above. In some embodiments, the vehicle braking rate and/or the maximum acceleration capability of the host vehicle may be determined based on, for example, the current speed of the vehicle, road conditions (e.g., a slope of the road, a material of the road, etc.), weather conditions (e.g., snowy, humid, etc.), vehicle conditions (e.g., tire pressure, brake-pad condition, current load of the vehicle, etc.), or the like. In some embodiments, one or more of the factors may be determined based on an output of one or more sensors. For example, a host vehicle may contain an accelerometer, which may provide an output to the processing device, and that output may include a current speed of the host vehicle and/or an acceleration capacity of the host vehicle. In some embodiments, the processing device may determine a current speed of the host vehicle and use the current speed to determine a current vehicle braking rate and/or a current acceleration capability. For example, the vehicle braking rate of a vehicle traveling at a first speed (e.g., 15 km/hr) may be significantly different from the vehicle braking rate of the same vehicle traveling at a second speed (e.g., 30 km/hr). In some embodiments, one or more of the vehicle braking rate, the maximum acceleration capability, and/or the speed of the host vehicle may be assumed using a predetermined value. For example, the processing device may be configured to assume that the host vehicle has a maximum braking capability corresponding with a predetermined value corresponding with an average (or the worst) maximum braking capability of the type of vehicle associated with the host vehicle. In some embodiments, each factor may be determined based on external conditions of the road or temporary characteristics of the vehicle. For example, the vehicle braking rate of the host vehicle may be determined based on a sensed condition of a road surface. In this example, the sensed road condition may include a roughness of the road, a slant or slope of the road, the present or absence of a substance or object on the road, whether the road is asphalt, cement, gravel, or another material, or any other condition consistent with this disclosure. As another example, the vehicle braking rate of the host vehicle may be determined based on a sensed weather condition. In this example, the weather condition may include a detection of any precipitation (e.g., rain, sleet, snow, ice, etc.), a weather condition that affects visibility (e.g., fog, smog, smoke, etc.), a weather condition that may affect the handling of the vehicle (e.g., strong winds, high heat, etc.), or any other weather condition that may affect a navigational response of the host vehicle. In another example, processing device may determine a vehicle braking rate based on whether the host vehicle contains, for example, one passenger or a plurality of passengers, cargo of a significant weight, a trailer, etc. Such conditions may be sensed based on analysis of images acquired by onboard cameras (e.g., recognizing snow or water on a roadway, water on a windshield, road surface smoothness, etc.) or may be based on any other suitable sensors. Such sensors may include wheel slip sensors, braking system sensor outputs, receivers of signals transmitted from roadside infrastructure, weather data streamed via a WiFi connection, GPS and/or speedometer outputs, etc.

[0799] In some embodiments, the vehicle braking rate and/or the maximum acceleration capability may be determined based on one or more predefined factors. For example, a government or industry entity may provide one or more regulations that dictate a vehicle braking rate and/or a maximum acceleration capability that a vehicle or class of vehicles may have and the at least one processor may
5 assume that that the host vehicle has the vehicle braking rate and/or the maximum acceleration capability allowed by the regulations. Constraints on the maximum braking and/or maximum acceleration values may also be user-defined (e.g., an occupant of an autonomous vehicle may select one or more vehicle setting that correspond to predetermined maximum allowable braking and acceleration levels), scenario-defined (e.g., based on circumstances at a particular location and/or during a particular navigation
10 scenario), and/or based on observed performance (e.g., based on historical data collected during autonomous navigation and/or navigation by a user).

[0800] Process 5000 may include a step 5012 for determining a stopping distance for the host vehicle. Consistent with this disclosure, the processing device may be configured to determine a stopping distance for the host vehicle based on the host vehicle braking rate, the host vehicle maximum
15 acceleration capability, and/or the current speed of the host vehicle (e.g., as described in detail above). The stopping distance for the host vehicle may be determined by any means consistent with this disclosure. For example, the processing device may use one or more of the formulas discussed in relation with RSS above.

[0801] In some embodiments, a stopping distance may be a distance that a host vehicle needs
20 to come to a stop, given its current speed, if the host vehicle were to accelerate at its maximum acceleration capability for period of time before braking at a predetermined braking capability. For example, the stopping distance for the host vehicle may include an acceleration distance that corresponds to a distance the host vehicle can travel over a predetermined time period at the current maximum acceleration capability of the host vehicle, starting from the determined current speed of the host vehicle.
25 The predetermined time period may be a time period assumed by the processing device. For example, a constraint may dictate that the processing device assumes a host vehicle will accelerate for a certain time period (e.g., 0.5 milliseconds, 1 millisecond, 2 milliseconds, 3 milliseconds, 10 milliseconds, etc.). The predetermined time period may be a reaction time associated with the host vehicle. For example, the predetermined time period may be a delay between the time when a sensor (e.g., an imaging device) of
30 the host vehicle captures an image of a condition that requires a braking response from the host vehicle (e.g., a target vehicle coming into the path of the host vehicle) and the time when the host vehicle begins braking at its vehicle braking rate. Such time periods may include processing times associated with analysis of the captured image, planning a navigational response based on one or more processor-
35 implemented driving policies, and implementing the navigational response through control of one or more actuators associated with a steering and/or braking system (among others) of a vehicle. In this example, it may be assumed that, in a worst-case scenario, the host vehicle accelerates at its maximum acceleration capability for the entirety of the delay period. In some embodiments, the time period associated with the reaction time may be determined or approximated by the processing device. For example, the processing

device may monitor the time between receiving an image and determining that a braking response is needed. As another example, the processing device may determine a mean, median, or mode reaction time associated with a plurality of navigational responses and use the mean, median, or mode reaction time when calculating a stopping distance. For example, after tens, hundreds, or thousands of navigational responses, the processing device may determine that the average time between a sensed event and a navigational response is a particular value (e.g., any real number of milliseconds) and use that value as the reaction time associated with the host vehicle. In some cases, the reaction time used in the RSS distance determination may be vehicle-specific and may include a running average of a measured reaction time for the particular vehicle to various sensed conditions.

[0802] Consistent with this disclosure, a determined stopping distance may include a minimum distance between the host vehicle and the target vehicle after coming to a complete stop. The determined or predetermined minimum distance may correspond with a predetermined separation distance to be maintained between the host vehicle and other vehicles. For example, the processing device may be programmed to include a determined or predetermined minimum distance in a calculation of a stopping distance. For example, a predetermined separation distance may be at least one meter. In this example, the determined or predetermined minimum distance may be d_{min} as discussed above.

[0803] Process 5000 may include a step 5014 for determining a current speed of the target vehicle, a target vehicle maximum acceleration capability, and a target vehicle braking rate. Consistent with this disclosure, the at least one processing device may be programmed to determine a current speed of a target vehicle. In some embodiments, the speed of the target vehicle may be determined based on the speed of the host vehicle. For example, one or more sensors on a host vehicle may provide information related to the speed of the host vehicle and an indication of the acceleration or change in direction of a target vehicle in relation to the host vehicle. In some embodiments, the target vehicle speed may be determined based on analysis of at least one image. The at least one image may be, for example, the same image received by the processing device at step 5002 that was used to identify the target vehicle at step 5006. In some embodiments, a plurality of images may be analyzed to determine a speed of the target vehicle. For example, a first image may depict the target vehicle a first distance from the host vehicle and a second image may depict the target vehicle a second distance from the host vehicle, the processing device may use the change in distance and the time between the first and second images to determine a speed associated with the target vehicle. In some embodiments, the target vehicle speed may be determined based on analysis of an output from at least one of a LIDAR system or a RADAR system associated with the host vehicle. For example, the processing device may use the speed of the host vehicle and a distance between the host vehicle and the target vehicle (as sensed by, for example, a LIDAR system) to determine a speed of the host vehicle.

[0804] In some embodiments, to account for a worst-case scenario, the processing device may be programmed to assume that the current speed of the target vehicle is less than or more than a sensed speed of the target vehicle. For example, if the target vehicle is traveling in front of the host vehicle, the processing device may be programmed to reduce a sensed speed using a predetermined value, a

percentage, a formula, or the like. If the processing device determined, based on information from one or more sensors, that a target vehicle is traveling at 100 km/hr, it may adjust the sensed speed by 10% to assume that the vehicle is traveling at 90 km/hr. The speed of a target vehicle may be correlated with the braking rate of the target vehicle so assuming a reduced speed of the target vehicle is akin to assuming the target vehicle may stop quicker than it otherwise might (e.g., a target vehicle may stop in a shorter distance if traveling at 90 km/hr than if traveling at 100 km/hr).

[0805] Consistent with this disclosure, the at least one processing device may be programmed to assume a maximum acceleration capability and/or vehicle braking rate of the target vehicle based on at least one recognized characteristic of the target vehicle. The at least one characteristic may be any characteristic that is known to correlate with an acceleration or braking capability or that may correlate with an acceleration or braking capability. In some embodiments, the recognized characteristic of the target vehicle may include a vehicle type, class, make, model, size, or the like, as described above.

[0806] The processing device may be programmed to recognize the at least one characteristic from information from one or more sensors (e.g., LIDAR, RADAR, imaging devices, etc.). In some embodiments, the recognized characteristic of the target vehicle may be determined based on analysis of the at least one image. For example, the image of the target vehicle may include a text on the exterior of the target vehicle, which may be used to determine a vehicle type, vehicle model, vehicle size, or other characteristic. As another example, the image of the target vehicle may include other objects that may be used to approximate a vehicle size or other characteristic based on comparison with the other objects. As an example, a large vehicle may appear taller than a speed-limit sign whereas a small vehicle may appear shorter than a speed-limit sign. In some embodiments, the at least one characteristic of the target vehicle may be determined based on at least one of a LIDAR output or a RADAR output. For example, a LIDAR output may include a first distance associated with the distance to the rear bumper of the target vehicle and a second distance associated with a distance to the front bumper (or other area) of the target vehicle, the difference between which may be used to estimate the size of the target vehicle. Sizes of target vehicles may also be estimated based on known sizes of imaged features (e.g., brake lights, tires, side view mirrors, etc.). A determination of how an object scales from one image to another may also be used to predict a target vehicle size or range.

[0807] As discussed above, the target vehicle braking rate may correspond to a maximal or submaximal braking rate of the vehicle. The target vehicle braking rate or target vehicle maximum acceleration capability may be determined by any means consistent with this disclosure, including the means described above in step 5010 for the host vehicle. For example, the target vehicle braking rate may be determined based on a sensed condition of a road surface, a sensed weather condition, etc. In some embodiments, the vehicle braking rate may depend on a responsibility or blame characteristic associated with the host vehicle or target vehicle. Accordingly, the host vehicle braking rate may be less than the target vehicle braking rate, or vice versa.

[0808] Process 5000 may include a step 5016 for determining a stopping distance for the target vehicle based on the target vehicle braking rate, the target vehicle maximum acceleration capability, and

the current speed of the target vehicle. The stopping distance for the target vehicle may be a distance that the target vehicle needs to come to a stop, given its current speed, if the target vehicle were to accelerate at its maximum acceleration capability for period of time before braking at a predetermined braking rate. For example, the stopping distance for the target vehicle may include an acceleration distance that
5 corresponds to a distance the target vehicle can travel over a predetermined time period at the current maximum acceleration capability of the target vehicle, starting from the determined current speed of the target vehicle. The distance may be determined by any means consistent with this disclosure, including the means described above in step 5012 with respect to the host vehicle.

[0809] Process 5000 may include a step 5018 for implementing the planned navigational
10 action if the planned navigation action is safe. For the purposes of step 5018, one example of when the planned navigation action may be considered safe is when if the determined next-state distance is greater than a sum of the stopping distance for the host vehicle and the stopping distance of the target vehicle. In this circumstance, there will be no collision between the host vehicle and the target vehicle because the stopping distance for both of the vehicles (including the potential distance traveled assuming maximum
15 acceleration during the reactionary period) would allow the vehicles to come to a complete stop before colliding if the navigational action was performed.

[0810] For the purposes of step 5018, another example of when the planned navigation action may be considered safe is when the determined next-state distance is greater, by at least a predetermined minimum distance, than the sum of the stopping distance for the host vehicle and the stopping distance of
20 the target vehicle. In this embodiment, the predetermined minimum distance may be d_{min} and may be predetermined, or determined, by any means consistent with this disclosure and have any attribute discussed above. In some embodiments, the predetermined minimum distance may correspond to a predetermined separation distance to be maintained between the host vehicle and other vehicles. For example, a constraint of an autonomous navigation system may dictate that the autonomous vehicle never
25 come within (less than) a specified distance of other vehicles. The predetermined separation distance may be a distance of any dimension. In some embodiments, the predetermined separation distance is at least one meter, two meters, five meters, etc. In some embodiments, the predetermined distance may vary depending upon, for example, the speed at which the host vehicle is traveling, the location of the host vehicle (e.g., an urban road, a rural highway, etc.).

30 [0811] Navigation with a Safe Lateral Distance

[0812] As described above, RSS may be effective for determining a safe distance between a host vehicle and a target vehicle in a longitudinal direction. In other embodiments, however, RSS may also be used to maintain a safe lateral distance between two vehicles. For example, a host vehicle and a target vehicle may be traveling relative to one another such that they are spaced laterally. The disclosed
35 embodiments may be used to ensure that a host vehicle maintains a safe lateral distance with the target vehicle in view of one or more navigational actions that may result in a change in lateral separation between a host and target vehicle. Maintenance of a safe lateral spacing may be especially important when a host vehicle and a target vehicle are driving next to one another. In such cases, unchecked lateral

travel by one or more of the vehicles could result in a collision. Safe lateral spacing, however, may also be of concern even where a host vehicle and target vehicle are longitudinally spaced such that a continuous lateral motion by one or more of the host vehicle and the target vehicle would not result in a collision. For example, where a lateral motion of one of the vehicles may cause a cut in situation such that an unsafe longitudinal distance would result following a lateral motion by one or more of the host vehicle and target vehicle, then there may be an interest in maintaining a certain safe lateral distance. The following sections describe various aspects of lateral navigation in terms of lateral spacing/separation, lateral acceleration, lateral velocity, and/or lateral braking. Such terms may refer to changes in a lateral component of motion of a vehicle. For example, in a situation where a host vehicle and a target vehicle are traveling in the same direction in adjacent lanes, if the host vehicle steers in a manner that causes the host vehicle to approach the target vehicle, it may be said that the lateral spacing between the host and target vehicle decreases. This decrease in lateral spacing may occur at a particular lateral velocity, which may be constant or which may change according to a lateral acceleration value. Lateral velocities associated with a decreasing lateral spacing may be designated with one sign (e.g., -), and lateral velocities associated with an increasing lateral spacing may be designated with an opposite sign (e.g., +). Lateral acceleration may refer to situations in which a host vehicle is laterally approaching a target vehicle (or laterally moving away from a target vehicle) at an increasing lateral velocity. The term “lateral braking” may refer to situations in which a host vehicle laterally approaches a target vehicle, but with a decreasing lateral velocity. For example, a maneuver resulting in lateral braking would cause the rate at which two cars are approaching one another to slow. Eventually, the lateral braking may result in a lateral velocity between the cars of 0. Maneuvers affecting lateral spacing, velocity, acceleration, etc. may include steering of one or more vehicles away from or toward another vehicle.

[0813] Fig. 51A illustrates an example scenario with two vehicles spaced laterally from each other consistent with the disclosed embodiments. In this scenario, a host vehicle 5102 and a target vehicle 5104 may be traveling alongside each other on a road, each in their own respective lanes. As shown in Fig. 51A, host vehicle 5102 may perform a maneuver, such as a lane change maneuver, which reduces a lateral distance between host vehicle 5102 and target vehicle 5104. The scenario shown in Fig. 51A is by way of example only, and various other scenarios may arise in which host vehicle 5102 maintains a safe lateral distance between the target vehicle. For example, while host vehicle 5102 and target vehicle 5104 are shown to travel in the same general direction in Fig. 51A, the vehicles may also be traveling in opposite directions with some lateral space between the vehicles. In other instances, the vehicles may be traveling with some lateral spacing through an intersection, in a parking lot, in a merging pattern, or in other possible scenarios. The present disclosure is not limited to any particular scenario in which a host vehicle and target vehicle are traveling with some lateral spacing between them.

[0814] As shown in Fig. 51A, target vehicle 5104 may be traveling at a velocity represented by velocity vector 5124. Similarly, host vehicle 5102 may be traveling at a velocity represented by velocity vector 5123. The lateral velocity 5123 of host vehicle 5102 may be defined as the lateral component of velocity vector 5122, as shown in Fig. 51A. The lateral component may be defined in

various ways. For example, the lateral component may use a frame of reference based on the previous heading direction of host vehicle 5102, the current heading direction of target vehicle 5104, the lane of travel of host vehicle 5102 and/or target vehicle 5104, or the like. In some embodiments, lateral velocity 5123 may be the component of the velocity of host vehicle 5102 that is in the direction of target vehicle 5104. The lateral velocity of target vehicle 5104 may be similarly defined as the component of a velocity of target vehicle 5104 that is in the direction of host vehicle 5102. Any other suitable means may be used to define the relative lateral velocity of the vehicles. Lateral spacing, lateral velocity, and/or lateral acceleration may be expressed in Cartesian coordinates, polar coordinates, or according to any other suitable convention.

[0815] Host vehicle 5102 may be configured to perform a planned navigational action if it allows host vehicle 5102 to maintain a safe longitudinal distance from vehicle 5104, such as an RSS safe distance, as discussed above. The planned navigational action may be an action that decreases the lateral distance between host vehicle 5102 and target vehicle 5104 (e.g., by increasing lateral velocity 5123). In the scenario shown in Fig. 51A, host vehicle 5102 and target vehicle 5104 may be traveling in parallel lanes with the same direction of travel and host vehicle 5102 may be changing lanes into the lane occupied by target vehicle 5104. Accordingly, the planned navigational action may include a lane change maneuver or merge maneuver. In other instances, the planned navigational action may include a passing maneuver, in which host vehicle 5102 is passing target vehicle 5104 and thus may temporarily move into an adjacent lane (either in the same or opposing direction of travel) and may move back into the lane occupied by target vehicle 5104. In other instances, the planned navigational action may include a turning maneuver. For example, host vehicle 5102 may be turning at an intersection, which may reduce the lateral distance between the vehicles. The planned navigational action may include any other maneuvers which may reduce a lateral distance between host vehicle 5102 and target vehicle 5104, such as exiting a highway, maneuvering around an object (e.g., an obstacle in the road, a pothole, a person, etc.), a parking maneuver (e.g., parallel parking, pulling into a parking spot, etc.), or other similar actions. In some instances, it may be target vehicle 5104 that performs the maneuver that reduces the lateral distance between target vehicle 5104 and host vehicle 5102. The planned navigational action of host vehicle 5102 may include a maintain forward heading maneuver, maintain throttle action, a maintain coasting action, change heading direction, braking maneuver (lateral and/or longitudinal), or the like.

[0816] As discussed above, the RSS safe distance may be the distance between a host vehicle and a target vehicle at which the host vehicle may come to a stop without colliding with the target. An analog to the longitudinal RSS safe distance may include an RSS safe lateral distance between the two vehicles. The safe lateral distance may include a calculation of the minimum distance that it would take the vehicles to stop laterally (e.g., reduce their lateral speed to zero). In this case, rather than determining the vehicle's braking capacity (e.g., applying the brakes), the safe lateral distance may be based on the host vehicle's ability to steer in a direction away from the other vehicle (or to otherwise cause a lateral approach velocity of two vehicles to slow to zero). Thus, the lateral braking distance may depend, at least in part, on the steering/turning capabilities of the vehicle.

[0817] Fig. 51B illustrates an example of a host vehicle 5102 maintaining a safe lateral distance consistent with the disclosed embodiments. Target vehicle 5104 may be traveling alongside host vehicle 5102 such that the vehicles are spaced laterally from each other. Similar to the longitudinal safe distance described above, a processing device of host vehicle 5102 (e.g., processing unit 110) may determine a planned navigational action for accomplishing a navigational goal of the host vehicle. Host vehicle 5102 may be configured to perform the planned navigational action if an RSS safe longitudinal distance can be maintained between host vehicle 5102 and target vehicle 5104. Additionally, host vehicle 5102 may be configured to perform the planned navigational action if an RSS safe lateral distance can be maintained between host vehicle 5102 and target vehicle 5104. As discussed above, the planned navigational action can include various maneuvers, such as a lane change maneuver, a merge maneuver, a passing maneuver, a parking lot navigation maneuver, a turning maneuver, an avoidance maneuver, a maintain forward heading maneuver, maintain throttle maneuver, or the like. The processing device of host vehicle 5102 may determine a next-state lateral distance 5110 between host vehicle 5102 and target vehicle 5104 corresponding to a lateral distance between the vehicles that would result if the planned navigational action was taken. The processing device may also determine a lateral braking distance 5112 for host vehicle 5102 and a lateral braking distance 5114 for target vehicle 5104. As noted above, lateral braking distances 5112 and 5114 may represent the lateral distance required for the vehicles to reduce their lateral velocity to zero. Accordingly, the vehicles may still have a nonzero overall velocity even upon laterally stopping. For example, the vehicles may continue traveling at some non-zero longitudinal velocity. The processing device of host vehicle 5102 may implement the planned navigational action if the determined next-state lateral distance 5110 is greater than a sum of lateral braking distance 5112 and lateral braking distance 5114, thereby indicating the vehicles can both reduce their lateral velocity (e.g., by turning) to avoid a collision. In the scenario shown in Fig. 51B, the navigational action of host vehicle 5102 would be allowed as the next-state lateral distance 5110 is greater than the sum of lateral braking distance 5112 and lateral braking distance 5114.

[0818] Consistent with the disclosed embodiments, lateral braking distance 5112 for host vehicle 5102 may correspond to a distance required to achieve a host vehicle lateral speed of zero. Lateral braking distance 5112 may depend on at a maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle, starting from the current lateral speed of the host vehicle. Similar to the maximum acceleration rate used in the longitudinal safe distance determination, described above, lateral braking distance 5112 may assume a lateral acceleration of host vehicle 5102. Accordingly, determination of lateral braking distance 5112 may assume a lateral velocity of host vehicle 5102 after a lateral acceleration of the host vehicle as a result of a turning maneuver made toward the target vehicle at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle during a reaction time associated with the host vehicle. Constraints on the maximum yaw rate capability and/or maximum change in turn radius capability may also be user-defined (e.g., an occupant of an autonomous vehicle may select one or more vehicle a setting that correspond to predetermined maximum allowable yaw rate and change in turn radius levels),

scenario-defined (e.g., based on circumstances at a particular location and/or during a particular navigation scenario), and/or based on observed performance (e.g., based on historical data collected during autonomous navigation and/or navigation by a user).

[0819] Fig. 52A illustrates an exemplary host vehicle 5102 during a planned navigation action
5 consistent with the disclosed embodiments. The yaw rate of host vehicle 5102 may represent the rate of change in a heading angle of host vehicle 5102, for example represented as an angular velocity, caused by turning the wheels. Host vehicle 5102 may have a heading direction 5205 at time t_1 and may have a heading direction 5206 at time t_2 . Heading direction 5205 may vary from heading direction 5206 by an angle, Θ . The yaw rate may be represented as $\frac{\Theta}{(t_2 - t_1)}$. In some instances, the yaw rate may be constant if
10 the turn radius is constant. For example, host vehicle 5102 may be turning at a constant turn radius 5207 and thus the yaw rate of host vehicle 5102 will remain constant. The change in turn radius of host vehicle 5102 may represent the rate in which host vehicle 5102 changes the turning radius. For example, host vehicle 5102 may decrease the turning radius to turn radius 5208, thereby increasing the yaw rate of host vehicle 5102. Host vehicle 5102 may also have a maximum yaw rate capability, which may correspond to
15 properties of the vehicle (e.g., the turn radius when the wheel is turned all the way to one side, etc.)

[0820] Similar to the longitudinal safe distance determination previously discussed, an RSS distance (e.g., a safe RSS lateral distance) be dependent upon the maximum possible lateral acceleration capability that the vehicle may experience during a period between an event and a vehicle's reaction to sensing of the event (e.g., a reaction time associated with the host vehicle). Accordingly, the lateral
20 braking distance may account not only for the lateral distance needed to slow the vehicle to a lateral speed of zero, but also a worst-case scenario of maximum lateral acceleration over a reaction time. The lateral braking distance may include a lateral acceleration distance that corresponds to a lateral distance the vehicle can travel during a vehicle reaction time at the vehicle's maximum yaw rate capability and a maximum change in turn radius capability for the vehicle. For example, the RSS lateral distance may
25 account for a delay before host vehicle 5102 or target vehicle 5104 begins performing a corrective steering maneuver. There may be several milliseconds (or more or less) between the time host vehicle 5102 detects that target vehicle 5104 is moving laterally towards host vehicle 5102 or vice versa. The processing device may account for the change in turn radius and yaw rate that may occur during response period ρ , corresponding to the vehicle's reaction time. Response period ρ may correspond to a known
30 characteristic of a vehicle, similar to the braking rate, or may be an assumed value for all cars. The lateral speed used in determining the lateral braking distance of host vehicle 5102 may assume a maximum yaw rate capability of host vehicle 5102 and that host vehicle 5102 increases the yaw rate based on a maximum change in turn radius capability of the vehicle during response period ρ , thereby representing a worst-case lateral braking distance 5112.

[0821] Referring again to Fig. 51B, the lateral braking distance 5114 for target vehicle 5104
35 may be calculated in a similar manner as lateral braking distance 5112 for target vehicle 5102. Lateral braking distance 5114 for target vehicle 5104 may correspond to a distance required to achieve a target vehicle lateral speed of zero. Lateral braking distance 5114 may depend on a maximum yaw rate

capability of the target vehicle and the maximum change in turn radius capability of the target vehicle, starting from the current lateral speed of the target vehicle. Lateral braking distance 5114 may assume a lateral acceleration of target vehicle 5104 and, accordingly, the determination of lateral braking distance 5114 may assume a lateral speed of host vehicle 5104 after a lateral acceleration of the host vehicle as a result of a turning maneuver by the target vehicle toward the host vehicle made at the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability during a reaction time associated with the target vehicle.

[0822] In some embodiments, the lateral speed of the host vehicle and/or the target vehicle may be determined based on an output of one or more sensors. For example, the lateral speed may be determined from an output of an accelerator, a LIDAR system, a RADAR system, a GPS unit, or the like. As another example, the lateral speed of the host vehicle may be determined by analyzing one or more images (e.g., based on a scaling rate change of a fixed object detected in two or more images). Similarly, the lateral speed of the target vehicle may be determined by analyzing one or more images (e.g., based on a scaling rate change of the target vehicle, etc.).

[0823] For a target vehicle, the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability may be determined based on at least one characteristic of the target vehicle, similar to the braking capability described above with respect to the longitudinal safe driving distance. The at least one characteristic may be any characteristic that may correlate with a yaw rate or turn radius capability of a vehicle. In some embodiments, the recognized characteristic of the target vehicle may include a vehicle type (e.g., full-size sedan, compact car, SUV, cross-over SUV, motorcycle, etc.), a vehicle make or model, brand name, or other classifier of the target vehicle (e.g., Toyota Prius®, BMW X5®, etc.), a vehicle size (as described above with respect to the longitudinal safe driving distance determination), or the like. The processing device may be programmed to assume the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability based on known values of that vehicle type or a class of vehicles. The processing device may use the vehicle characteristic, for example, to access a database containing known maximum yaw rate capabilities and/or maximum change in turn radius capabilities associated with each vehicle type, model, etc. The characteristic of the vehicle may be determined or identified through analysis of an image.

[0824] For host vehicle 5102, the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle may be known or estimated by the system based on characteristics of the vehicle, data from previous turning maneuvers, or the like. In some embodiments, the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle may be based on predetermined constraints for the host vehicle. For example, the maximum lateral acceleration may be less than the vehicle's true maximum capabilities but may be limited by a setting. A maximum lateral braking rate or maximum lateral acceleration can be uncomfortable to passengers, can excessively wear down vehicle hardware (tires, brake pads etc.), or the like. Accordingly, the predetermined constraints (which may be set by system 100, a driver of host vehicle 5102, etc.) may limit the maximum yaw rate capability of the host vehicle and the maximum

change in turn radius capability of the host vehicle. These limited values may be used in the determination of the lateral braking distance.

[0825] In some embodiments, at least one of the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, the target vehicle maximum yaw rate capability, or the target vehicle maximum change in turn radius capability may be determined based on one or more sensed road conditions (e.g., road material type, road traction properties, etc.), one or more sensed weather conditions (e.g., rain, snow, fog, ice), or the like. For example, a processing device may determine maximum yaw rate capability and/or the maximum change in turn radius capability that provides the maximum safety and comfort on a snowy road, which may differ from those that would be determined for a dry highway, and so forth.

[0826] In some embodiments the lateral braking rate may be based on a blame or responsibility determination associated with host vehicle 5102 or target vehicle 5104. From a responsibility perspective, host vehicle 5102 may have an obligation to steer away from target vehicle 5104, but at a lower rate than target vehicle 5104. For example, target vehicle 5104 may be veering into the lane of host vehicle 5102, and thus may have a greater obligation to steer away from host vehicle 5102. Accordingly, the host vehicle maximum yaw rate capability and/or maximum change in turn radius capability may be less than those of the target vehicle, or vice versa, even for the same vehicle type or model.

[0827] In some embodiments, the processing device of host vehicle 5102 may be configured to implement the planned navigational action if the determined next-state distance is greater, by at least a predetermined minimum distance, than the sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle. If the RSS safe longitudinal distance includes this minimum approach distance component, then when the host vehicle 5102 reaches a zero lateral speed, the host vehicle should remain separated from the target vehicle by at least the minimum approach distance. The minimum distance may correspond to a predetermined lateral separation distance (e.g., a minimum distance of 0.5 m, 1 m, 2 m, etc.) to be maintained between the host vehicle and other vehicles. The predetermined lateral separation may be determined based on any set of factors consistent with this disclosure (e.g., a velocity associated with host vehicle 5102 and/or target vehicle 5104, a sensed weather condition, a user's preferences, etc.).

[0828] In some embodiments, the host vehicle may also account for the longitudinal distance between the host vehicle and the target vehicle. Fig. 52B shows an example in which host vehicle 5202 may perform a navigation action based on both a lateral and longitudinal safe driving distance. Host vehicle 5202 may determine whether to perform a planned navigation action based on a lateral distance between host vehicle 5202 and target vehicle 5204. The planned navigation action may be any of the maneuvers described above, such as a lane change maneuver, a merge maneuver, etc. The processing device of host vehicle 5202 may determine a next-state lateral distance 5210, as well as a lateral braking distance 5212 for the host vehicle and a target vehicle lateral braking distance 5214. In the scenario shown in Fig. 52B, the sum of the lateral braking distance 5212 for the host vehicle and target vehicle

lateral braking distance 5214 may be greater than next-state lateral distance 5210, indicating that the navigation action would not result in a safe lateral distance from target vehicle 5204. As a result, the vehicle may forgo the planned navigational action that would result in an unsafe lateral distance.

[0829] In some embodiments, however, the processing device may be programmed to
5 implement the planned navigational action regardless of the next-state lateral distance (e.g., if the planned navigational action would result in a safe longitudinal distance 5220 between the host vehicle and the target vehicle). In other words, in cases where a navigational action can be taken, and the next state distance would be greater than or equal to the safe RSS longitudinal distance, then whether the lateral separation between a host vehicle and a target vehicle is greater than or equal to a safe RSS lateral
10 distance may be less relevant (e.g., the host vehicle may be significantly forward of or behind a target vehicle, such that there is little or no risk of a lateral collision). The safe longitudinal distance may be determined in any manner consistent with the present disclosure, including the systems and methods described in detail above. For example, safe longitudinal distance 5220 (which may correspond to the next-state longitudinal distance described above) may be greater than or equal to a sum of a host vehicle
15 longitudinal stopping distance 5222 and a target vehicle longitudinal stopping distance 5224. Accordingly, in the scenario shown in Fig. 52B, the processing device of host vehicle 5202 may perform the planned navigation action. As discussed above, the host vehicle stopping distance may be based on a maximum longitudinal braking capability of the host vehicle and a current longitudinal speed of the host vehicle. The host vehicle stopping distance may further include a host vehicle acceleration distance
20 corresponding to a distance the host vehicle can travel at a maximum acceleration capability of the host vehicle during a reaction time associated with the host vehicle starting from the current longitudinal speed of the host vehicle. Similarly, the target vehicle stopping distance may be based on a maximum longitudinal braking capability of the target vehicle and a current longitudinal speed of the target vehicle. The target vehicle stopping distance may include a target vehicle acceleration distance corresponding to a
25 distance the target vehicle can travel at a maximum acceleration capability of the target vehicle during a reaction time associated with the target vehicle starting from the current longitudinal speed of the target vehicle.

[0830] Figs. 53A and 53B include flowcharts depicting an exemplary process 5300 for
maintaining a safe lateral distance consistent with the disclosed embodiments. Process 5300 may be
30 performed by at least one processing device, such as processing unit 110, described above. Process 5300 is exemplary only. One of ordinary skill in the art having the benefit of this disclosure may understand that process 5300 may include additional steps, exclude certain steps, or may be otherwise modified in a manner consistent with this disclosure.

[0831] Process 5300 may include a step 5302 for receiving at least one image representative of
35 an environment of the host vehicle. The at least one image may be received from an image capture device, which may be any image capture device consistent with this disclosure, including image capture device 122. In some embodiments, the at least one image may be an image obtained from any of a camera, a RADAR, a LIDAR, or any other device from which an image may be obtained, whether optical

or otherwise. There may be some period of time between when the image is captured and when the processing device receives the image. Further, there may be a period of delay between the time when an event occurs and the imaging device captures an image of the event. For example, if a target vehicle enters a lane in front of a host vehicle, there may be a short period (e.g., a millisecond, 2 milliseconds, 5 milliseconds, etc.) between when the target vehicle maneuvers into the lane and when the imaging device captures an image of the target vehicle. There may also be time associated with analysis of the capture image, development of a planned response based on the analysis, transmission of one or more control signals to actuation systems, and an actual response of the actuation systems to the control signal(s).

[0832] Process 5300 may include a step 5304 for determining a planned navigational action for the host vehicle. Consistent with this disclosure, the at least one processing device may be programmed to determine a planned navigation action for accomplishing a navigational goal of the host vehicle. The navigation action may be determined based on at least one driving policy. The planned navigation action and/or the at least one driving policy may be any consistent with this disclosure, including those discussed above. For example, the planned navigation action may include at least one of a lane change maneuver, a merge maneuver, a passing maneuver, a parking lot navigation maneuver, or a turning maneuver.

[0833] Process 5300 may include a step 5306 for analyzing the at least one image to identify a target vehicle in the environment of the host vehicle. The at least one image may be an image received from an image capture device, such as image capture device 122. The at least one image may be the one or more images received as part of step 5302 of process 5300.

[0834] Process 5300 may include a step 5308 for determining a next-state lateral distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken. Consistent with this disclosure, the processing device may be programmed to determine the next-state distance between the host vehicle and the target vehicle. The next-state lateral distance may be calculated by any means disclosed herein. For example, if the planned navigation action is a lane change maneuver of the host vehicle, the next-state distance may be a distance between the host vehicle and a target vehicle in a lane next to the host vehicle as a result of the navigational action.

[0835] Process 5300 may include a step 5310 for determining a maximum yaw rate capability of the host vehicle, a maximum change in turn radius capability of the host vehicle, and a current lateral speed of the host vehicle. Each of these factors (i.e., the maximum yaw rate capability, the maximum change in turn radius capability, and the current lateral speed of the host vehicle) may be determined by any means consistent with this disclosure. The maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle may depend on the properties of the host vehicle (e.g., the turn radius of the host vehicle, tire pressure, tire wear, etc.). In some embodiments, the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle are based on predetermined constraints. For example, the maximum yaw rate capability of the host vehicle and/or the maximum change in turn radius capability of the host vehicle may be constrained based on a user or system setting (e.g., to prevent uncomfortable lateral accelerations, prevent

unnecessary wear or damage to the vehicle, etc.). Accordingly, the maximum yaw rate capability of the host vehicle and/or the maximum change in turn radius capability of the host vehicle may depend on the current speed of the host vehicle. For example, the maximum yaw rate capability of a vehicle traveling at a first speed (e.g., 15 km/hr) may be significantly different from the maximum yaw rate capability of the same vehicle traveling at a second speed (e.g., 30 km/hr). In some embodiments, at least one of the maximum yaw rate capability of the host vehicle or the maximum change in turn radius capability of the host vehicle is determined based on a condition of a road surface (e.g., a slope of the road, a material of the road, the presence of a substance on the road, etc.), weather conditions (e.g., snowy, humid, etc.), vehicle conditions (e.g., tire pressure, brake-pad condition, current load of the vehicle, etc.), or the like. In some embodiments, one or more of the factors may be determined based on an output of one or more sensors. In another example, processing device may determine a maximum yaw rate capability and/or maximum change in turn radius capability based on whether the host vehicle contains, for example, one passenger or a plurality of passengers, cargo of a significant weight, a trailer, etc.

[0836] In some embodiments, the maximum yaw rate capability and/or maximum change in turn radius capability may be determined based on one or more predefined factors. For example, a government or industry entity may provide one or more regulations that dictate a maximum yaw rate capability and/or maximum change in turn radius capability that a vehicle or class of vehicles may have and the at least one processor may assume that that the host vehicle has the vehicle braking rate and/or the maximum acceleration capability allowed by the regulations.

[0837] Process 5300 may include a step 5312 for determining a lateral braking distance for the host vehicle. Consistent with this disclosure, the processing device may be configured to determine a lateral braking distance for the host vehicle based on the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, and the current lateral speed of the host vehicle. The lateral braking distance for the host vehicle may be determined by any means consistent with this disclosure.

[0838] In some embodiments, the lateral braking distance for the host vehicle may correspond to a distance required to achieve a host vehicle lateral speed of zero. The lateral braking distance may be the distance required to reach a lateral speed of zero at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle, starting from the current lateral speed of the host vehicle and after a lateral acceleration of the host vehicle as a result of a turning maneuver made toward the target vehicle at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle during a reaction time associated with the host vehicle. For example, the lateral braking distance for the host vehicle may include a lateral acceleration distance that corresponds to a distance the host vehicle can travel over a predetermined time period assuming the host vehicle turns at the maximum rate towards the target vehicle. The predetermined time period may be a time period assumed by the processing device. For example, a constraint may dictate that the processing device assumes a host vehicle will accelerate for a certain time period (e.g., 0.5 milliseconds, 1 millisecond, 2 milliseconds, 3 milliseconds, 10 milliseconds, etc.). The predetermined

time period may be a reaction time associated with the host vehicle. For example, the predetermined time period may be a delay between the time when a sensor (e.g., an imaging device) of the host vehicle senses a condition that requires a lateral braking response from the host vehicle (e.g., a target vehicle coming into the path of the host vehicle) and the time when the host vehicle begins braking at its lateral braking rate. In this example, it may be assumed that, in a worst-case scenario, the host vehicle accelerates laterally at its maximum capability for the entirety of the delay period. In some embodiments, the time period associated with the reaction time may be determined or approximated by the processing device. For example, the processing device may monitor the time between receiving an image and determining that a lateral movement response is needed. As another example, the processing device may determine a mean, median, or mode reaction time associated with a plurality of navigational responses and use the mean, median, or mode reaction time when calculating a lateral braking distance. For example, after tens, hundreds, or thousands of navigational responses, the processing device may determine that the average time between a sensed event and a navigational response is a particular value (e.g., any real number of milliseconds) and use that value as the reaction time associated with the host vehicle.

[0839] Process 5300 may include a step 5314 for determining a current lateral speed of the target vehicle, a target vehicle maximum yaw rate capability, and a target vehicle maximum change in turn radius capability. Consistent with this disclosure, the at least one processing device may be programmed to determine a current lateral speed of a target vehicle. In some embodiments, the lateral speed of the target vehicle may be determined based on the lateral speed of the host vehicle. For example, one or more sensors on a host vehicle may provide information related to the speed of the host vehicle and an indication of the acceleration or change in direction of a target vehicle in relation to the host vehicle. In some embodiments, the target vehicle lateral speed may be determined based on analysis of at least one image. The at least one image may be, for example, the same image received by the processing device at step 5302 that was used to identify the target vehicle at step 5306. In some embodiments, a plurality of images may be analyzed to determine a lateral speed of the target vehicle. For example, a first image may depict the target vehicle a first distance from the host vehicle and a second image may depict the target vehicle a second distance from the host vehicle, the processing device may use the change in distance and the time between the first and second images to determine a lateral speed associated with the target vehicle. In some embodiments, the target vehicle lateral speed may be determined based on analysis of an output from at least one of a LIDAR system or a RADAR system associated with the host vehicle. For example, the processing device may use the lateral speed of the host vehicle and a distance between the host vehicle and the target vehicle (as sensed by, for example, a LIDAR system) to determine a lateral speed of the host vehicle.

[0840] In some embodiments, to account for a worst-case scenario, the processing device may be programmed to assume that the current lateral speed of the target vehicle is less than or more than a sensed lateral speed of the target vehicle. For example, if the target vehicle is traveling next to the host vehicle, the processing device may be programmed to reduce a sensed speed using a predetermined value, a percentage, a formula, or the like.

[0841] Consistent with this disclosure, the at least one processing device may be programmed to assume a maximum yaw rate capability and/or a maximum change in turn radius capability of the target vehicle based on at least one recognized characteristic of the target vehicle. The at least one characteristic may be any characteristic that may correlate with a maximum yaw rate or turn radius capability. In some embodiments, the recognized characteristic of the target vehicle may include a vehicle type, class, make, model, size, or the like, as described above.

[0842] The processing device may be programmed to recognize the at least one characteristic from information from one or more sensors (e.g., LIDAR, RADAR, imaging devices, etc.). In some embodiments, the recognized characteristic of the target vehicle may be determined based on analysis of the at least one image. For example, the image of the target vehicle may include a text on the exterior of the target vehicle, which may be used to determine a vehicle type, vehicle model, vehicle size, or other characteristic. As another example, the image of the target vehicle may include other objects that may be used to approximate a vehicle size or other characteristic based on comparison with the other objects. As a particular example, a large vehicle may appear taller than a speed-limit sign whereas a small vehicle may appear shorter than a speed-limit sign. In some embodiments, the at least one characteristic of the target vehicle may be determined based on at least one of a LIDAR output or a RADAR output. For example, a LIDAR output may include a first distance associated with the distance to the rear bumper of the target vehicle and a second distance associated with a distance to the front bumper (or other area) of the target vehicle, the difference between which may be used to estimate the size of the target vehicle.

[0843] As discussed above, the target vehicle braking rate may correspond to a maximal or submaximal braking rate of the vehicle. The target vehicle braking rate or target vehicle maximum acceleration capability may be determined by any means consistent with this disclosure, including the means described above in step 5310 for the host vehicle. For example, the target vehicle braking rate may be determined based on a sensed condition of a road surface, a sensed weather condition, etc. In some embodiments, the vehicle braking rate may depend on a responsibility or blame characteristic associated with the host vehicle or target vehicle. Accordingly, the host vehicle braking rate may be less than the target vehicle braking rate, or vice versa.

[0844] Process 5300 may include a step 5316 for determining a lateral braking distance for the target vehicle based on the current lateral speed of the target vehicle, the target vehicle maximum yaw rate capability, and the target vehicle maximum change in turn radius capability. The lateral braking distance for the target vehicle may correspond to a distance required to achieve a target vehicle lateral speed of zero. The lateral braking distance may be the distance required to reach a lateral speed of zero at the maximum yaw rate capability of the target vehicle and the maximum change in turn radius capability of the target vehicle, starting from the current lateral speed of the target vehicle and after a lateral acceleration of the target vehicle as a result of a turning maneuver by the target vehicle toward the host vehicle made at the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability during a reaction time associated with the target vehicle. For example, the lateral braking for the target vehicle may include a lateral acceleration distance that corresponds to a distance the

target vehicle can travel over a predetermined time period assuming the target vehicle turns at the maximum rate towards the host vehicle. The predetermined time period may be a time period assumed by the processing device to be a reaction time associated with the target vehicle. The reaction time for target vehicle may be the same as the reaction time for the host vehicle or may be different. For example, the processing device may assume a reaction time based on a human driver, which may be significantly different than the reaction time associated with an automated vehicle.

[0845] Process 5300 may include a step 5318 for implementing the planned navigational action if the planned navigation action is safe. For the purposes of step 5318, one example of when the planned navigation action may be considered safe is when the planned navigational action if the determined next-state lateral distance is greater than a sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle. In this circumstance, there will be no collision between the host vehicle and the target vehicle because the lateral braking distance for both of the vehicles (including the potential distance traveled assuming maximum lateral acceleration during the reactionary period) would allow the vehicles to steer away from each other before colliding if the navigational action was performed.

[0846] In some embodiments, the processing device may be configured to implement the planned navigational action if the determined next-state lateral distance is greater, by at least a predetermined minimum distance, than the sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle. The predetermined minimum distance may correspond to a predetermined lateral separation distance to be maintained between the host vehicle and other vehicles. For example, a constraint of an autonomous navigation system may dictate that the autonomous vehicle never come within (less than) a specified distance of other vehicles. The predetermined separation distance may be a distance of any dimension. In some embodiments, the predetermined separation distance is at least 1 m, 2 m, 5 m, etc. In some embodiments, the predetermined distance may vary depending upon, for example, the speed at which the host vehicle is traveling, the location of the host vehicle (e.g., an urban road, a rural highway, etc.).

[0847] In other embodiments, the processing device may be programmed to implement the planned navigational action regardless of the next-state lateral distance if the planned navigational action would result in a safe longitudinal distance between the host vehicle and the target vehicle, as described above with respect to Fig. 52B. For example, the safe longitudinal distance may be greater than or equal to a sum of a host vehicle longitudinal stopping distance and a target vehicle longitudinal stopping distance. The host vehicle stopping distance may be based on a maximum longitudinal braking capability of the host vehicle and a current longitudinal speed of the host vehicle, and the host vehicle stopping distance may include a host vehicle acceleration distance corresponding to a distance the host vehicle can travel at a maximum acceleration capability of the host vehicle during a reaction time associated with the host vehicle starting from the current longitudinal speed of the host vehicle. The target vehicle stopping distance may be based on a maximum longitudinal braking capability of the target vehicle and a current longitudinal speed of the target vehicle, and the target vehicle stopping distance may include a target

vehicle acceleration distance corresponding to a distance the target vehicle can travel at a maximum acceleration capability of the target vehicle during a reaction time associated with the target vehicle starting from the current longitudinal speed of the target vehicle.

[0848] Navigation Relative to Pedestrians at Crosswalks

5 [0849] The present disclosure describes a system (also referred to as a navigation system) for navigating a host vehicle in proximity to a pedestrian crosswalk. The navigation system may be configured to identify traffic lights along a roadway traveled by an autonomous or semi-autonomous host vehicle. The navigation system may be configured to receive information from various vehicles about locations of various traffic lights along the roadway, map the locations of the traffic lights on a sparse
10 map available to the system and to the vehicles, and receive from the vehicles information related to the traffic lights, as well as information related to pedestrians, and crosswalks. For example, when an autonomous vehicle approaches a crosswalk, the system may be configured to receive information about the state of a relevant traffic light (e.g., the traffic light includes a green light), information whether pedestrians are present at the crosswalk, locations of pedestrians, their motion, their predicted future
15 motion, and the like.

[0850] The navigation system may include multiple computing devices, processors, data storage devices, databases, and the like for efficient information processing. The navigation system may include non-transitory storage devices or computer-readable media. In some embodiments, the storage devices may include hard drives, compact discs, flash memory, magnetic-based memory devices, optical-based
20 memory devices, and the like. The navigation system may include one or more processors configured to perform instructions that may be stored on one or more non-transitory storage devices associated with the navigation system. The navigation system may include a distributed navigation system and vehicle navigation systems. The distributed navigation system may include multiple components that may not be connected to any particular vehicle. The components may communicate via a network, which in some
25 cases, may be a 5G network. The distributed navigation system may include multiple processing computing systems located in the vicinity of various roadways to reduce the network latency time associated with information transmission between the distributed navigation system and vehicle navigation systems. The vehicle navigation system may be part of a vehicle control system associated with a vehicle. In various embodiments, the vehicle navigation systems and the distributed navigation
30 system may exchange information related to any aspects of navigation of autonomous, semi-autonomous and/or human-operated vehicles via any suitable wireless communication (e.g., WiFi, 4G, 5G, communication based on visible light, infrared, microwaves, radio waves, and the like). The vehicle navigation systems and the distributed navigation system form the overall navigation system.

[0851] The navigation system (e.g., vehicle navigation systems and distributed navigation
35 system) may determine a navigational action for the host vehicle that may include steering, braking, or accelerating the host vehicle. The navigational action may include parameters such as rate of steering, rate of braking, or rate of acceleration. The navigational action may include actions that may not be directly related to the motion of a vehicle. For example, such navigational actions may include turning on/off

headlights, engaging/disengaging antilock brakes, switching transmission gears, adjusting parameters of a vehicle suspension, turning on/off vehicle warning lights, turning on/off vehicle turning lights or brake lights, producing audible signals and the like. In various embodiments, the navigational actions are based on navigational data available to the navigation system.

5 [0852] The navigational information available to the navigation system may include any suitable information about various vehicles, traffic lights, traffic signals, and pedestrians in the proximity of the host vehicle. Examples of navigational information may include the position of various autonomous and human-operated vehicles in proximity of the host vehicle, velocities of the vehicles, accelerations of the vehicles, lengths, heights, widths, and masses of the vehicles, accelerating and braking abilities of the
10 vehicles, predicted trajectories of the vehicles, navigational actions executed by the vehicles, locations of traffic lights, states of traffic lights, locations of crosswalks, locations of traffic signs (including roadway surface signs), locations of pedestrians, movements of pedestrians, and the like.

[0853] In various embodiments, the navigation system may receive, from an image capture device, at least one image representative of an environment of the host vehicle. The image capture device
15 may be a camera for capturing visible light. In some cases, however, the image capture device may be an infrared camera, an ultraviolet camera, microwave detecting device, LIDARs, RADARs, and the like. In various embodiments, the image capture device may include one or more cameras. For example, the image capture device may include multiple cameras that are pointing in different directions. The images captured by the image capture device may represent the environment in front of the host vehicle, behind
20 the host vehicle, at the left side of the host vehicle, at the right side of the host vehicle, the environment of a roadway in front of the host vehicle, and the like. Further, the image capture device may capture images of traffic lights along the roadway, images of pedestrians in proximity of roadway, images of crosswalks, and the like. In some cases, the image capture device may be movable (e.g., rotatable) to capture images of the environment of the host vehicle from at a selected azimuth and polar angle. In various
25 embodiments, images collected from image capture devices from different vehicles may be transmitted to a distributed navigation system for further analysis. The collected images may also be distributed to different vehicle navigation systems.

[0854] The navigation system may detect, based on an analysis of the at least one image, captured by the image capture device, a representation of a pedestrian crosswalk in the at least one image.
30 In an example embodiment, the crosswalk may be recognized in an image using any suitable recognition technique (e.g., machine image recognition technique based on neural networks). For instance, the crosswalk may be recognized by a characteristic marking of a roadway (e.g., white rectangular regions painted on a surface of the roadway across the roadway). In some cases, crosswalk may be identified in
an image based on additional information about the location of the host vehicle. For example, the
35 navigation system may determine the location of the host vehicle and identify a crosswalk in proximity to the location of the host vehicle. For example, the navigation system may maintain information about locations of crosswalks. In an example embodiment, the navigation system may determine the location of the host vehicle, the location of the crosswalk and orient the image capture device to correctly capture the

image of the crosswalk. In some cases, based on the location of the host vehicle, and location of various other vehicles presented in the vicinity of the crosswalk, the navigation system may be configured to capture images of the crosswalk from different angles using image capturing devices of different vehicles. The captured images then may be distributed to the vehicles in the vicinity of the crosswalk.

5 [0855] In some embodiments, image capture devices may be installed at the crosswalk. For example, the image capture devices may be part of the traffic lights, traffic signals, and the like. These devices may also be used by the navigation system to capture images of the crosswalk and transmit images to a host vehicle navigation system located in the vicinity of the crosswalk.

10 [0856] In various embodiments, the navigation system may analyze one or more images of the crosswalk and determine, based on analysis of the images, whether a pedestrian is present at (or in the vicinity) of the crosswalk. Pedestrians may be identified using any suitable image recognition technology (e.g., machine-learning techniques). In an example embodiment, pedestrians may be identified by image capture devices installed at a crosswalk by comparing images of a crosswalk taken at different times. For example, if at a first time an image does not contain a representation of a pedestrian, but at a second time
15 an image contains a representation of a pedestrian, the navigation system may determine that a pedestrian is near a crosswalk. Furthermore, multiple images of a crosswalk may improve identification of a pedestrian at a crosswalk. For example, multiple images of a pedestrian approaching the crosswalk may help identify the pedestrian at a crosswalk. In various embodiments, the navigation system may determine if a representation of a pedestrian appears in the at least one image captured by one or more image
20 capturing devices.

[0857] Besides analyzing images of a crosswalk to identify pedestrians, the navigation system may detect, based on analysis of the at least one image, a presence of a traffic light in the environment of the host vehicle. The traffic lights may be detected in the environment of the host vehicle (e.g., near a crosswalk, to the left or to the right of the host vehicle, and the like). The navigation system may detect
25 the presence of the traffic light in the environment of the host vehicle using any suitable approaches which may not necessarily be based on analysis of images of the environment of the host vehicle. For example, the detection of the presence of the traffic light in the environment of the host vehicle may be based on one or more signals transmitted by the detected traffic light. In some cases, however, a state of a detected traffic light may be determined based on analysis of captured images. For example, a color
30 associated with an illuminated lamp may be detected to determine whether a traffic light is in a stop state (e.g., red light detected), in a go state (e.g., green light detected), etc. The location of lamps within a traffic light housing can also be used to determine or to indicate a state of a traffic light, as in some jurisdictions, a go state may be associated with illumination of a lowest vertically positioned lamp and a stop state may be associated with a highest vertically positioned lamp.

35 [0858] In an example embodiment, the traffic light may communicate with an autonomous vehicle using wireless communication to enable determination of a state of the traffic light (or relevancy of the traffic light to a particular host vehicle or to a lane of travel of the host vehicle). The wireless communication may include one or more devices configured to exchange transmissions over an air

interface to one or more networks (e.g., cellular, the Internet, etc.) using an electromagnetic field at the radio frequency, infrared frequency, or ultraviolet frequency. Additionally, or alternatively, wireless communication may use magnetic fields, electric fields, or sound. Such transmissions can include communications between a traffic light and an autonomous vehicle in the proximity of the traffic light, and/or in some cases, such communication may include communication between the traffic light and the distributed navigation system and between the distributed navigation system and the vehicle navigation system of an autonomous vehicle. It should be noted, that various other forms of communication between the traffic light and the vehicle may be used that may include audio communication (e.g., communication using soundwaves such as ultrasound, laser-based communications and the like). In some cases, communications may include time-dependent visible communications (e.g., the time-dependent signal produced by LED sources).

[0859] In various embodiments, a traffic light may include a traffic light communication system for generating a signal to an autonomous vehicle. For example, the traffic light communication system may include a signal generating module for generating a “stop” signal, a “go” signal, a “slowdown” signal, a “speedup” signal and the like. For example, the “stop” signal may indicate that the autonomous vehicle needs to stop abruptly. Such signal, for example, may correspond to a red light. The “go” signal may indicate that the vehicle needs to start or continue moving. Such signal may correspond, for example, to a green light. The “slowdown” signal may indicate that the autonomous vehicle needs to slow down. In some embodiments, the “slowdown” signal may include additional parameters incorporated in information related to the signal that may include the required speed for the vehicle or deceleration of the vehicle. In an example embodiment, the “slowdown” signal may correspond to yellow light. The “speed up” signal may indicate that the autonomous vehicle needs to increase its speed. In an example embodiment, the traffic light communication system may transmit to an autonomous vehicle the duration of time the traffic light maintains a particular state, and/or the remaining time during which the traffic light will maintain a particular state. For example, if the traffic light is in a green state, it may transmit to the autonomous vehicle that the state will be switched to a yellow state in thirty seconds.

[0860] In some cases, one or more autonomous vehicles may interact with the traffic light and request state timing changes for the traffic light. For example, when there are no pedestrians at the crosswalk of a roadway, autonomous vehicles may request a traffic light to maintain its state by transmitting the request to the traffic light via wireless communication. The navigation system may process various requests from autonomous vehicles and adjust the state timing of the traffic light. It should be noted that various signals described above are illustrative and various other signals may be incorporated as well. For example, traffic light generated signals may indicate that the vehicle needs to turn to the right, turn to the left, change lanes, or make a U turn.

[0861] In various embodiments, in addition to generating signals, the traffic light communication system may receive navigational information from various vehicles and generate signals based on the received information. For example, the traffic light communication system may receive a host vehicle’s

speed and a distance of the host vehicle from the traffic light and may generate deceleration request for the host vehicle based on the host vehicle's speed and the distance from the traffic light.

[0862] In various embodiments, the traffic light wireless communication may include various ways to authenticate communication from/to a traffic light as well as to provide secure communication between the traffic light and an autonomous vehicle. Additionally, or alternatively, secure communication may be established between the autonomous vehicle and the distributed navigation system. In an example embodiment, secure communication may be established through the use of private and public keys. For example, the autonomous vehicle and the distributed navigation system may exchange the public key for encrypting the secure information and may use private keys for information decryption. Similarly, the autonomous vehicle may exchange public keys with the traffic light communication system. In some embodiments, the traffic light communication system may authenticate through the distributed navigation system. For example, the traffic light communication system may provide identification and/or password information to the server of the distributed navigation system, and the server may issue a secure token to the traffic light communication system. In some embodiment, the server may encrypt the secure token using the public key of the traffic light communication system and transmit the encrypted token to the traffic light communication system. Further, the server may be configured to encrypt the secure token using public key associated with an autonomous vehicle and transmit the encrypted token to the autonomous vehicle. In various embodiments, the traffic communication system may include the secure token for a communication data packet transmitted to the autonomous vehicle to provide authentication for the vehicle. Additionally, or alternatively, the traffic communication system may include the secure token for a communication data packet transmitted to the server. In a similar way, the autonomous vehicle may include the secure token for a communication data packet transmitted to the traffic light communication system. Additionally, or alternatively, the autonomous vehicle may include the secure token in a communication data packet transmitted to the server. It should be noted that the secure communication process described above is illustrative, and various other approaches may be used. The authenticated secure communication between various autonomous vehicles, traffic lights, and the distributed navigation system (as well as secure communication among various autonomous vehicles) may ensure that system for navigation of various autonomous vehicles cannot be compromised by a third party (e.g., a party attempting to alter movements of the autonomous vehicles).

[0863] In various embodiments, the navigation system may determine whether the detected traffic light is relevant to the host vehicle and the pedestrian crosswalk. The determination of the relevancy of the traffic light may be based on analysis of images collected by the navigation system. For example, from a captured image, a particular traffic light may be identified above a certain lane of travel of a roadway represented in the captured image. Such information can enable the processor(s) of the host vehicle to determine whether the detected traffic light is or may be relevant to the host vehicle.

[0864] Additionally, or alternatively, the relevancy of a traffic light may be determined using navigational information available to the navigation system (e.g., locations of the traffic lights, locations of crosswalks, location of the host vehicle, a vector from the host vehicle to neighboring traffic lights, and

the like). In an example embodiment, the relevancy of the traffic light to the crosswalk and the host vehicle may be determined based on traffic light relevancy information included in one of maps accessible by the navigation system. For example, through crowd sourcing, identifiers associated with detected traffic lights along with their corresponding locations may be stored in one or more maps (e.g., REM maps). In addition to this information, traffic light relevancy to certain lanes of a roadway may also be stored in the maps. This information, for example, may be based on crowdsourced information gathered from previous drives where, for example, cars drove through an intersection using lanes 1, 2, and 3 when detected traffic lights X and Y were determined to be in a green state. Of course, there are many other examples of how vehicle behavior can be tracked and correlated to traffic light state to indicate traffic light relevancy to particular lanes of a roadway.

[0865] The relevancy of a traffic light (especially with respect to crosswalks detected in the vicinity of a traffic light) may be determined in a similar manner based on a detected state of the traffic light and motion of various other vehicles in the vicinity of the crosswalk. For example, if the state for a traffic light corresponds to a red light state and a vehicle stops in front of the traffic light and prior to entering the crosswalk, the navigation system may determine that the traffic light is relevant to the crosswalk and a lane traversed by the host vehicle. In some embodiments, the navigation system may be configured to receive various states corresponding to the traffic lights as well as navigational information from various autonomous vehicles in order to determine the relevancy of various traffic lights. For instance, a first vehicle may communicate to the navigation system a first state for a traffic light and first navigational information associated with the movement of the first vehicle, and a second vehicle may communicate to the navigation system a second state for the traffic light and a second navigational information associated with the movement of the second vehicle. This information may be incorporated into the pool of information used to construct the REM maps for a particular roadway, which can store traffic light locations along with their relevancy to lanes of travel, crosswalks, etc.

[0866] In various embodiments, a traffic light may be determined as relevant to a certain crosswalk. For example, based on image analysis a particular traffic light may be determined to be facing the crosswalk in a direction generally perpendicular to a lane of travel of a vehicle. In some cases, the detected traffic light may display special symbols (e.g., a symbol of a person walking, a symbol of a person standing, a symbol of a palm indicating a stop request, and the like). These features may be analyzed to identify the traffic light's relevancy to a crosswalk. This determination may be made in real time as a host vehicle navigates and intersection to assist the host vehicle in navigating relative to a detected crosswalk, or it may be analyzed in a post-capture process where drive information from multiple previous drives are collected, analyzed, and used to generate a map that indicates traffic light relevancy. Such maps may be used by a particular host vehicle to determine, for example, whether traffic light detected in an image is relevant to a current lane of travel and/or whether the detected traffic light is a pedestrian crosswalk signal relevant to pedestrians navigating a crosswalk detected in an captured image. It should be noted that various relevancy criteria for a traffic light discussed above are illustrative, and other relevancy criteria for the traffic light that utilizes the navigation system may be used. For instance,

in some cases, lanes of a roadway may be identified by numerical, alphanumerical values or any other suitable identifiers by the navigation system. The navigational system may identify the relevant traffic light for each lane. For such cases, the determination of relevancy of the traffic light to the host vehicle amounts to determining the lane for the host vehicle. In various embodiments, the navigation system may determine the lane identifier for the lane traversed by the host vehicle using, for example, the position of the host vehicle, distance to the left or right curbside of the roadway, and the like. Additionally, or alternatively, the lane may have markers related to its identifier positioned along the lane that may be detected and identified by the host vehicle navigation system.

[0867] Traffic light relevancy to a particular crosswalk may also be determined based on previous drive information in which pedestrians were detected in captured images walking through a particular crosswalk with a particular nearby traffic light was in a red state. Such information may enable a determination that when the traffic light is in a red state, pedestrians are free to enter the crosswalk. Such information may also suggest that when the same traffic light is in a green state (and cars are free to pass the traffic light), that pedestrians should not enter the crosswalk. This information can also be stored in the REM maps (e.g., as indicators of traffic light relevancy to nearby crosswalks) and used by vehicles to navigate relative to the mapped crosswalks and traffic lights.

[0868] As discussed above, a traffic light in the vicinity of the crosswalk may be configured to communicate with various vehicles using wireless communication. For instance, the traffic light may communicate to the host vehicle navigation system a lane identifier of a lane for which the traffic light is relevant. The host vehicle navigation system may compare the lane identifier received from a traffic light with the lane identifier obtained by the host vehicle via analysis of the navigational information and/or images of the crosswalk containing a representation of the traffic light. If the lane identifier received from the traffic light matches the lane identifier obtained by host vehicle via, for example, analysis of the images of the crosswalk or analysis of the navigational information, then traffic light may be determined to be relevant to the host vehicle and to the crosswalk.

[0869] In various embodiments, when it is determined that one or more pedestrians appear in at least one image in the vicinity of the crosswalk, the navigation system may determine a proximity of the pedestrian relative to the detected pedestrian crosswalk. For example, the proximity of the pedestrian may be determined by analyzing images of the crosswalk or a vicinity of the crosswalk. In some cases, the estimate of a distance between the pedestrian approaching the crosswalk and the curbside of the road (at which the crosswalk starts) may be improved by analyzing multiple images and determining an averaged distance obtained by averaging different distance estimates obtained from multiple images. In an alternative embodiment, the crosswalk may include various sensors for detecting the presence of pedestrians and the proximity of the pedestrians relative to the detected pedestrian crosswalk. Such sensors may include location sensors, RADAR, LIDAR, etc. The data acquired by sensors may be transmitted to the navigation system for further analysis and processing.

[0870] The navigation system may be configured to determine based on at least one driving policy, a planned navigational action for causing the host vehicle to navigate relative to the detected

pedestrian crosswalk. The driving policy may be, for example, a set of rules/constraints for driving an autonomous vehicle in the vicinity of a crosswalk (and elsewhere). For example, the driving policy may include guidelines for navigating through an intersection associated with a crosswalk (e.g., what speed to travel when navigating the intersection, how far to stay away from pedestrians and other vehicles, where to stop relative to the crosswalk should a stop be needed, etc.). The driving policy may also embody rules or guidelines for navigating relative to a detected crosswalk in view a state of a detected, relevant traffic light. For example, as noted above, a crosswalk and traffic light may be detected based on analysis of captured images (often in a single captured image). Relevancy of the detected traffic light to the detected crosswalk may be determined, for example, by accessing a REM map that stores traffic light locations (e.g., to confirm a mapped traffic light corresponds with the image-detected traffic light) and relevancy of traffic lights to lanes/crosswalks. Alternatively or additionally, the traffic light relevancy information may be determined based on relevancy information transmitted from the detected traffic light (or by any other suitable technique).

[0871] In view of a determination that a traffic light is relevant to a particular crosswalk, the driving policy may determine appropriate navigational actions for a host vehicle based on detected states of the traffic light and detected states of the crosswalk. For example, such navigational actions may include stopping at a predetermined distance before the crosswalk on a red traffic light, slowing down to a first velocity (or maintaining a velocity of the host vehicle at a first velocity) when traveling through the crosswalk when the traffic light is green and when there are no pedestrians detected in a vicinity of the crosswalk, and slowing down to a second velocity (e.g., slower than the first velocity) when the traffic light is green and there is at least one pedestrian detected in the vicinity of the crosswalk. It should be noted, that the driving policy may depend on one or more of the type of a crosswalk, the width of the roadway, the number of lanes on the roadway, the number of vehicles in front, behind, to the left or to the right of the host vehicle, whether the other vehicles are moving or stationary in the vicinity of the crosswalk, the estimated trajectories of the vehicles, the velocities of other vehicles, the type of vehicles (e.g., whether the vehicles in the vicinity of the crosswalk are large, small, slow, or fast), a presence of intersection, the number of roads at the intersection, the number of crosswalks at the intersection, and the like. Additionally, the driving policy may depend on one or more of a number of pedestrians at the crosswalk, locations of pedestrians at the crosswalk, the estimated movements of various pedestrians, the estimated trajectories of the pedestrians, and the like.

[0872] The determination of the navigational actions in the proximity of a crosswalk may be based on various factors, regardless of the driving policy. In some embodiments, the navigational actions may be determined using a rule-based computer model, which may include a set of instructions implemented as software code. Additionally, or alternatively, the navigational actions may be determined using a machine-learning models, such as neural networks, generative-adversarial neural networks (GANs), information maximizing GANs (InfoGANs), and the like. The machine learning models may be trained using training data collected from various sensors (e.g., image sensors for collecting images of an environment of an example vehicle, sensors for collecting audio data relate to the environment of the

vehicle, and the like. Additionally, the training data may include navigational actions executed by human drivers, or by trained machine-learning modes that correspond to the images of the environment of the example vehicle. In an example embodiment, the navigational actions for the example vehicle may be deduced by analyzing data of various sensors of the example vehicle, such as acceleration sensors, sensors for collecting information about position of a steering wheel of a vehicle, sensors for collecting information about breaking, sensors for collecting information about vehicle transmission gears, and the like). The machine-learning model then may be trained to replicate the navigational actions executed by the human drivers base on an input of the images and/or audio of the environment of the example vehicle. In an example embodiment, the visual and audio data collected by sensors may be input to the machine-learning model, such as a neural network, and the neural network may return as an output the acceleration data for a vehicle and the wheel orientation of the vehicle, which may then be used to generate one or more the driving instructions for the vehicle.

[0873] In various embodiments, the determination of the planned navigational action may be based on the determined state of the detected traffic light and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk. In an example embodiment, if the detected traffic light is green, and there are no pedestrians detected in a proximity of the detected crosswalk, the navigation system may determine a navigational action that includes traveling through the crosswalk at a reduced speed, where the vehicle speed may depend on conditions associated with an environment of the vehicle (e.g., weather conditions around the vehicle, time of the day, visibility around the vehicle, whether the crosswalk is located at a busy intersection, the presence of vehicles in the vicinity of the crosswalk, and the like). In other words, if the traffic light state is green and there are no pedestrians detected near the crosswalk, it may be deduced that the host vehicle is unlikely to encounter a pedestrian entering the crosswalk before or during a time when the host vehicle transits the crosswalk. In cases where the detected traffic light is green, and there is a pedestrian detected crossing the crosswalk, the vehicle may determine the navigational action that includes slowing and stopping the vehicle. That is, even if the light is green, suggesting that the host vehicle has the right of way, the host vehicle should be navigated in a manner such that there is no risk of a collision with the pedestrian detected in the crosswalk or entering the crosswalk. In other cases, when there is a green light detected, but a pedestrian is also detected in proximity to the crosswalk (e.g., within a certain predetermined distance, which may be associated with a distance the pedestrian could potentially travel toward the crosswalk in a time that the host vehicle takes to reach and/or traverse the crosswalk), the navigation system may determine to slow the vehicle to a safe speed, which may be selected to allow the vehicle to abruptly stop if the pedestrian is detected entering the. In some cases, the navigational actions may include moving a vehicle into a different lane of a roadway to create a space between the vehicle and a pedestrian detected approaching the crosswalk or standing in proximity to the crosswalk. In an example embodiment, when the detected light is red, the navigation system may determine a navigational action including slowing and stopping the vehicle prior to reaching the crosswalk. Thus, in determining a navigational action for causing the host vehicle to navigate relative to a detected crosswalk, the processor(s) of the host vehicle may rely upon both a state

of a detected traffic light (and whether the detected traffic light is determined to be relevant to the host vehicle lane of travel and/or to the detected crosswalk) and whether pedestrians are detected in a vicinity of a the detected crosswalk.

[0874] In various embodiments, after determining the navigational action, the navigation system may interact with a vehicle control system to cause one or more actuator systems of the host vehicle to implement the planned navigational action. In an example embodiments, the actuators for the host vehicle may include vehicle steering actuators (e.g., systems responsible for steering of the host vehicle), vehicle accelerating and braking actuators (e.g., systems responsible for braking or accelerating the host vehicle), sound actuators (e.g., systems responsible for making a sound), light actuators (e.g., systems responsible for flashing light), and the like.

[0875] An example of a roadway containing several crosswalks is illustrated in Fig. 54. Fig. 54 shows an intersection 5401, including crosswalks 5451, 5452, and 5453. Intersection 5401 includes multiple lanes such as lane 5403 traveled by a vehicle 5441, and lane 5404 traveled by a vehicle 5442. Vehicle 5441 travels along a roadway 5413 that may include a left curbside (a curbside 5407) and a right curbside (a curbside 5402). As shown in Fig. 54, intersection 5401 may have multiple traffic lights such as traffic lights 5430, 5432, and 5433. In an example embodiment, traffic light 5430 may include an image capture device 5460 for capturing images of intersection 5401. In some embodiments, capture device 5460 may have multiple cameras for capturing various elements of the intersection. For example, one camera of capture device 5460 may capture crosswalk 5451 and areas of the roadway in the vicinity of crosswalk 5451, a second camera of capture device 5460 may capture crosswalk 5452 and areas of the roadway in the vicinity of crosswalk 5452, a third camera may capture crosswalk 5453 and areas of the roadway in the vicinity of crosswalk 5453. Additionally, capture device 5460 may include cameras that may zoom to focus on pedestrians, such as pedestrians 5420, 5422, 5421, 5423, and 5424, take images of bicyclists (e.g., a bicyclist 5425, as shown in Fig. 54) traversing intersection 5401 as well as vehicles (e.g., vehicles 5441 and 5442, as shown in Fig. 54) approaching intersection 5401. In some embodiments, capture device 5460 may zoom to focus on facial expressions of the pedestrians and/or capture multiple images of the pedestrians to provide the navigation system with information to determine movements of the pedestrians, and the like. In various embodiments, vehicle 5441 may have an image capture device 5461 and vehicle 5442 may have a respective image capture device 5462 for capturing various images of intersection of 5401, as previously described. Image capture devices 5461 and 5642 may have multiple movable cameras and may capture images of different portions of the intersection, similar to image capture device 5460.

[0876] Fig. 54 shows that various pedestrians may have different abilities, velocities, expected trajectories, and the like. For example, pedestrian 5422 may be determined to enter crosswalk 5451 and traverse crosswalk 5451 before traffic light 5432 has a red light (i.e., the light prohibiting the pedestrians traversing crosswalk 5451). Pedestrian 5420, on the other hand, appears to be stationary and may be unlikely to enter crosswalk 5451. Pedestrian 5424 may be a young person on a skateboard, whose motion may be unpredictable. Pedestrian 5421 may be an old gentleman who traverses crosswalk 5453 slowly

and predictably. Similarly, pedestrian 5423 is a person in a wheelchair who has limited mobility (e.g., unlikely to rapidly change direction when traversing crosswalk 5453). Any or all of these detected characteristics may be relied upon in determining a navigational action for the host vehicle, including, for example, a velocity at which the host vehicle is allowed to approach or traverse a crosswalk.

5 [0877] In various embodiments, the navigation system may include a machine-learning model for evaluating age, gender, expected movement, predictability of motion for a pedestrian identified at the crosswalk. For example, the navigation system may include a machine-learning model that can take, as an input, images obtained by one or more image capture devices (e.g., devices 5460, 5461, and 5462), and return an age of a pedestrian, a motion vector of the pedestrian, a location of the pedestrian, an expected
10 trajectory of the pedestrian, an uncertainty trajectory envelop (e.g., a trajectory envelop 5471 for pedestrian 5424, or a trajectory envelop 5472 for pedestrian 5422), and the like. A trajectory envelope may indicate possible trajectories for a pedestrian. As can be seen in Fig. 54, trajectory envelop 5471 may be wide for pedestrian 5424 and covers crosswalk 5452 and 5453, while envelop 5472 may be narrow for pedestrian 5422 and covers crosswalk 5451. Such identified characteristics may also be used in
15 determining a navigational action for the host vehicle relative to a detected crosswalk (e.g., whether to stop the host vehicle prior to the crosswalk, whether to slow the host vehicle, what velocity is appropriate for the host vehicle, etc. in navigating relative to the crosswalk).

[0878] In an example embodiment, the machine-learning model may be trained by observing movements of various pedestrians of different ages, different sizes, different genders, different physical
20 abilities, and the like to establish the most probable trajectory of a given pedestrian. In addition to determining the most probable motion, the machine-learning method may also output the ranges of possible motions, the likelihood of a most probable motion, the likelihood of possible deviations from the most probable motion and the like. In an example embodiment, if navigation system observes an erratic movement of a pedestrian, it may conclude that further movement of the pedestrian may be erratic.
25 Alternatively, the navigational system may conclude that some of the pedestrians may have a consistent motion (e.g., pedestrian 5421).

[0879] The machine-learning model described above is only illustrative, and any other suitable approaches for predicting pedestrian motion may be employed. For example, a rule-based approach may be used, or a rule-based approach combined with the machine-learning approach may be used.

30 [0880] The navigation system may make an assessment of the likelihood of a pedestrian entering the crosswalk prior to an intersection of the crosswalk by the host vehicle. For example, based on the images captured by one or more image capture devices (e.g., image capture devices associated with the host vehicle, with other vehicles in the vicinity of the crosswalk, image capture devices installed in the vicinity of the crosswalk, and the like), the navigation system may determine if the pedestrian is
35 approaching the crosswalk or traveling away from the crosswalk. In some cases, the navigational system may determine possible trajectories for pedestrians, and determine a navigational action for the host vehicle that is statistically safe and allows the host vehicle to travel along roadway 5413. The navigational action that is statistically safe may include an action that is estimated to be safe 99% of the time if a

trajectory of any one of pedestrians at intersection 5401 is unpredictable. However, the navigational action may be determined to be significantly safer (e.g., safe 99.99999% of the time) if trajectories of pedestrians are predictable. It should be noted that the safety of the pedestrians depends on actions taken by the pedestrians, and roadways may not be guaranteed to be completely safe regardless of the actions of the pedestrians. The observed or predicted actions/trajectories of the pedestrian may be combined with a detected state of the traffic light in determining a navigational action for the host vehicle.

[0881] In various embodiments, the planned navigational action for the host vehicle may include traversing the pedestrian crosswalk without a pedestrian-induced adjustment (i.e., any navigational action determined to avoid a collision with a pedestrian or reduce a risk of collision with an pedestrian even in view of an unexpected action of the pedestrian against a state of a detected traffic light (e.g., pedestrian walking into a crosswalk when a vehicle relevant traffic light is in a green state—suggesting the pedestrian has a stop signal)) if the detected traffic light is in a green state, and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk indicates that the pedestrian is not likely to enter the crosswalk prior to an intersection of the host vehicle with the crosswalk. In an example embodiment, the host vehicle (e.g., vehicle 5441, as shown in Fig. 54) may identify the state of traffic light 5430 as it approaches crosswalk 5451. If the state of traffic light 5430 is green and if no pedestrians are approaching crosswalk 5451, vehicle 5441 may not execute a pedestrian-induced adjustment. In various cases, the navigational system may determine whether to traverse crosswalk 5451 and how to traverse crosswalk 5451 (e.g., at what speed, at what acceleration, in which lane of roadway 5413, along which trajectory, with or without the sound of a horn or flashing of headlamps, and the like). Both decisions (i.e., whether to traverse crosswalk 5451 and how to traverse crosswalk 5451) may depend on the observed state of the traffic light, how long the traffic light is in the observed state, as well as the proximity of pedestrians to crosswalk 5451. For example, if a pedestrian is detected to be far from crosswalk 5451, vehicle 5441 may proceed according to the state of traffic light 5430 only (e.g., green state allowing for vehicle 5441 to traverse crosswalk 5451 and red state requiring vehicle 5441 to stop before reaching crosswalk 5451). In an example embodiment, if the detected pedestrian is close to crosswalk 5451, suggesting the pedestrian could step into crosswalk 5441 (e.g., pedestrian 5422 attempting to enter crosswalk 5441) and traffic light 5430 is in a green state, vehicle 5441 may proceed through the crosswalk, but may do so by taking at least some cautious action relative to pedestrian 5422. For example, vehicle 5441 may slow down, move away from pedestrian 5422, and the like, to minimize the risk of collision with pedestrian 5422 even if pedestrian 5422 disobeys a pedestrian walk indicator and enters the crosswalk. In various embodiments, if traffic light 5430 is in a green state and pedestrian 5422 enters any part of the crosswalk, vehicle 5441 may be required to stop, regardless of the state of traffic light 5430.

[0882] The navigation system may be configured to execute a planned navigational action that allows the host vehicle to traverse the pedestrian crosswalk (e.g., crosswalk 5451) without a pedestrian-induced adjustment if the detected traffic light (e.g., traffic light 5430) is in a green state and the determined proximity of a pedestrian relative to the detected pedestrian crosswalk, in combination with a

detected heading and speed associated with the pedestrian, indicates that the pedestrian cannot enter crosswalk 5430 prior to an intersection of the host vehicle with the crosswalk. For example, if a predicted trajectory of the pedestrian does not intersect the crosswalk, the host vehicle may proceed to traverse crosswalk (e.g., where the pedestrian is detected close to the crosswalk, but walking along a sidewalk perpendicular to the crosswalk). It should be noted, that the navigation system may evaluate position and trajectories of any or all the pedestrians in the proximity of crosswalk 5451 in order to determine that none of the pedestrians have possible trajectories that may intersect the trajectory of vehicle 5441.

[0883] If the navigation system determines that at least one of the pedestrians (e.g., pedestrian 5422, as shown in Fig. 54) may enter crosswalk 5451 prior to an intersection of vehicle 5441 with crosswalk 5451, the navigation system may include a planned navigational action for traversing crosswalk 5451 with a pedestrian-induced adjustment, if traffic light 5430 is in a green state.

[0884] In an example shown in Fig. 55A, if it is determined that pedestrian 5521 may enter crosswalk 5451, vehicle 5441 may determine a planned navigational action (e.g., a navigational action resulting in a trajectory 5502) for traversing crosswalk 5451 with a pedestrian-induced adjustment. In an example embodiment, the pedestrian-induced adjustment may include deflecting trajectory of vehicle 5441 as indicated by trajectories 5502 and 5503. For example, the pedestrian-adjustment of the navigational action resulting in trajectory 5503 may be due to fast-moving pedestrian 5521, and the adjustment resulting in trajectory 5502 may be due to slower moving pedestrian 5521. In various embodiments, in addition to monitoring pedestrians in the vicinity of crosswalk 5451, the navigational system may monitor locations and velocities of various vehicles (e.g., vehicle 5511) traveling in the vicinity of crosswalk 5451. In an example embodiment, if vehicle 5511 travels fast, the navigational system may choose trajectory 5502, and not trajectory 5503, to avoid collision with vehicle 5503. In some embodiments, a pedestrian-adjustment may include speeding up vehicle 5441, slowing or stopping vehicle 5441. For example, if pedestrian 5521 has already entered crosswalk 5451, vehicle 5441 may take trajectory 5504 that includes slowing down and stopping vehicle 5441, regardless of the state of a traffic light (e.g., traffic light 5430, as shown in Fig. 54). In some cases, another possible pedestrian-adjustment may include vehicle 5441 to increase its speed and choose trajectory 5502 or 5503 to pass pedestrian 5521. A pedestrian-induced adjustment may also include a reduction in host vehicle speed. The magnitude of the reduction in speed may be dependent upon a detected velocity of a pedestrian (e.g., the faster the pedestrian moves, the more the velocity of the host vehicle may be reduced), a predicted time the pedestrian may or could enter the crosswalk (the shorter the time, the more the host vehicle speed may be reduced), or based on any other detected characteristic of pedestrian movement relative to the crosswalk.

[0885] In some cases, the navigation system may execute multiple pedestrian-adjustments for multiple vehicles traveling along a roadway. For example, Fig. 55B shows vehicles 5441 and 5511 traveling along roadway 5413 towards crosswalk 5451, while a vehicle 5512 may travel in an opposite direction, as indicated by an arrow 5531. For cases, when it is determined that pedestrian 5521 may enter crosswalk 5451, a first pedestrian-adjustment for the navigational action for vehicle 5441 may result in

trajectory 5524, a second pedestrian-adjustment for the navigational action for vehicle 5511 may result in trajectory 5525, and a third pedestrian-adjustment for the navigational action for vehicle 5512 may result in trajectory 5526 (e.g., vehicle 5512 steering towards a curbside 5532 and stopping). In various embodiments, the pedestrian-induced adjustment may be determined based on a detected heading and speed of the pedestrian. For example, the heading and speed of the pedestrian may be determined taking multiple images of the pedestrian based on one or more image capture devices (e.g., devices 5460-5462, as shown in Fig. 54). In an example embodiment, as shown, for example, in Fig. 55A, the pedestrian-induced adjustment may include a lateral movement of vehicle 5441 away from the pedestrian (e.g., movement across roadway 5413 away from pedestrian 5521, as shown by trajectories 5502 and 5503).

Additionally, or alternatively, the pedestrian-induced adjustment may include vehicle 5441 slowing down prior to the intersection of vehicle 5441 with crosswalk 5451. Additionally, or alternatively, the navigational system may require vehicle 5441 to stop prior to the pedestrian crosswalk, even when traffic light (e.g., traffic light 5430, as shown in Fig. 54) is in a green state if pedestrian 5521 is detected to enter the pedestrian crosswalk.

[0886] In various embodiments, the distance of a pedestrian to a crosswalk may be determined using geometrical calculations. For example, Figs. 55C and Figs. 55D shows examples of calculating the distance of a pedestrian to a crosswalk using image capture devices capable of motion detection. In an example embodiment, shown in Fig. 55C, pedestrian 5421 approaches a starting location 5567 of crosswalk 5541. Image capture device 5460 may include a movable camera 5565 that may include a scope 5570 for accurate orientation of camera 5565 at a point in the proximity of pedestrian 5421 (e.g., a point 5571 next to shoes of pedestrian 5421, as shown in Fig. 55C). In some cases, movable camera 5565 may include a laser 5564 that may, for example, operate using an invisible wavelength (e.g., deep blue wavelength, infrared wavelength, microwave wavelength, and the like) and a detector for detecting a reflection spot of the laser beam from a surface (wherein the reflection spot is referred to as a laser marker). The movable camera may then be oriented by placing the laser marker in the proximity of pedestrian 5421 (e.g., at point 5571). Once the camera is oriented relative to starting location 5567 of crosswalk 5541, angles A , and B , as shown in Fig. 55C, may be determined. As shown in Fig. 55C, image capture device 5460 is mounted on a pole 5563 and is positioned above the surface of a crosswalk at a height H . For parameters A , B , and H , a distance 5561 of pedestrian 5421 to starting location 5567 of crosswalk 5451 is determined to be $H \cdot \tan(A) \cdot \sin(B)$.

[0887] Fig. 55D shows that a similar calculation may be made using image capture device 5461 associated with vehicle 5441. For device 5461 located at a height H_c above the surface of a roadway, with angles A_c and B_c determining the orientation of device 5461, the distance 5561 of pedestrian 5421 to starting location 5567 of crosswalk 5451 is determined to be $H_c \cdot \tan(A_c) \cdot \sin(B_c)$. In various embodiments, the navigation system may obtain an estimate of distance 5561 using different image capture devices and reduce the error in calculation using distance averaging, elimination of outlier estimates, and the like.

[0888] As described above, the navigational system may include multiple computing devices for executing various operations for a host vehicle navigation. An example process 5600 of navigating the host vehicle in the vicinity of the crosswalk is illustrated in Fig. 56. At step 5601 of process 5600, a processing device of the navigation system may receive a set of images representing an environment of the host vehicle. The images may be taken by one or more image capture devices as described above. The images may include the representation of pedestrians in the vicinity of a crosswalk, images of traffic lights, traffic signs, images of various vehicles in the vicinity of a crosswalk, and the like.

[0889] At step 5603 of process 5600, the processing device may detect a starting and an ending location of a crosswalk. In an example embodiment, the crosswalk, and the starting and the ending locations of the crosswalk may be detected by analyzing images taken by the one or more image capture devices. Alternatively, a location of the crosswalk may be identified based on a map associated with the navigational system. For example, the map may contain all the relevant information about crosswalk, and the navigation system may be configured to determine a positioning vector from the host vehicle to the center of the crosswalk, a positioning vector from the host vehicle to the starting location of the crosswalk and a positioning vector from the host vehicle to the ending location of the crosswalk. In case of the host vehicle approaching an intersection containing more than one crosswalk, multiple crosswalks may be identified on the map available to the navigation system, with respective starting and ending locations for each of the identified crosswalk.

[0890] At step 5605 of process 5600, the processing device may determine, based on an analysis of at least one image obtained by one or more image capture devices (e.g., image capture devices 5460-5462, as shown in Fig. 54) whether a pedestrian is present in the proximity of the crosswalk. In an example embodiment, as shown for example in Figs 55C and 55D the orientation of an image capture device (e.g., device 5460) may allow accurate determination of distance 5561 of pedestrian 5421 to start location of crosswalk 5567.

[0891] At step 5607 of process 5600, the processing device may detect the presence of a traffic light in the environment of the host vehicle. In an example embodiment, the processing device may detect the presence of the traffic light using at least one image of the environment of the vehicle. For instance, the processing device may receive an image of a crosswalk (e.g., crosswalk 5451, as shown in Fig. 54), and detect, based on a suitable image recognition algorithm, a traffic light (e.g., traffic light 5430, as shown in Fig. 54).

[0892] At step 5609 of process 5600, the processing device may determine a traffic light relevancy (e.g., the relevancy of traffic light 5430, as shown in Fig. 54) for a vehicle (e.g., vehicle 5441, as shown in Fig. 54). In an example embodiment, the relevancy of traffic light 5430 may be determined by analyzing a correlation between the time-dependent navigational information for various vehicles traveling along the road segment containing traffic light 5430 and the time-dependent state identifier for traffic light 5430 as described above. For example, if there is a direct correlation between the navigational information (or one or more changes in the navigational information) and the state identifier for traffic light 5430 (or changes in the state identifier), the processing device may determine that traffic light 5430

is relevant to vehicle 5441, to a lane (e.g. lane 5403, as shown in Fig. 54) traveled by the vehicle, and to crosswalk 5451, as shown in Fig. 54. In various embodiments, the correlation between the time-dependent navigational information for various vehicles traveling along the road segment containing traffic light 5430 and the time-dependent state identifier for traffic light 5430 may be obtained by
5 collecting information for multiple vehicles traveling along the road segment at different times.

[0893] Additionally, or alternatively, if there is a direct correlation between the movement of pedestrians and the state identifier for traffic light 5430 (or changes in the state identifier), the processing device may determine that traffic light 5430 is relevant to crosswalk 5451. For example, when traffic light 5430 is in a state red (as observed by vehicle 5441), and pedestrians are traversing the crosswalk, and
10 when traffic light 5430 is in a state green (as observed by vehicle 5441), and pedestrians are waiting at starting location 5567 (as shown in Fig. 55C) for the state of traffic light 5430 to change, the navigation system may conclude that traffic light 5430 is relevant to crosswalk 5451. In various embodiments, the correlation between the time-dependent movements of various pedestrians traversing crosswalk 5451 containing traffic light 5430 and the time-dependent state identifier for traffic light 5430 may be obtained
15 by collecting information for multiple pedestrians traversing crosswalk 5451 at different times.

[0894] In some embodiments, the relevance of traffic light 5430 to crosswalk 5451 and lane 5403 may be determined by referring to a map associated with the navigation system, where the relevance of traffic light 5430 may be known/stored/represented. For instance, one or more administrators of the navigation system may assign an identification number (ID number) to various traffic lights for a segment
20 of a roadway and associate their ID numbers and their locations with various crosswalks and lanes for the roadway segment.

[0895] At a step 5611 of process 5600, the processing device may determine a state of traffic light 5430 based on one or more images obtained by image capture devices (e.g., devices 5460-5462, as shown in Fig. 54). The processing device may compare the images obtained for various traffic lights with
25 images of various traffic lights (e.g., images of traffic lights having a green light, red light, yellow light, etc. stored in a memory of the navigation system) to determine states of various traffic lights.

[0896] Additionally, or alternatively, as described above, the state of traffic light 5430 may be communicated to a navigational system by traffic light 5430 using a wireless communication system. In various embodiments, a processing device may determine how much time is left until a change of the
30 state of traffic light 5430. In an example embodiment, the information about the time left until the change of the state of traffic light 5430 may be transmitted wirelessly by traffic light 5430, and in other embodiments, the information may be obtained by navigation system by collecting information of states of traffic light as a function of time by analyzing images of various vehicles traveling in the proximity of crosswalk 5430.

[0897] At step 5613 of process 5600, the processing device may determine a navigational action
35 for the host vehicle in vicinity of the detected crosswalk, based on the relevancy of traffic light 5430, the determined state of the traffic light, a presence of the pedestrian in the proximity of the crosswalk, a shortest distance selected between the pedestrian and one of the starting or the ending location of the

crosswalk, and a motion vector of the pedestrian. As described above, the navigational action may include slowing and stopping vehicle 5441 or moving vehicle laterally away from a pedestrian entering crosswalk 5451, or near crosswalk 5451.

[0898] At step 5615, if the navigational action is to be implemented, the processing device may cause one or more actuator systems associated with vehicle 5441 to implement the determined one or more navigational actions. In various embodiments, the one or more actuator systems may include regular controls for vehicle 5441 such as a gas pedal, a braking pedal, a transmission shifter, a steering wheel, a hand brake and the like. In some embodiments, actuator systems may be internal systems not accessible by a human operator that perform similar functions as the regular controls accessible to the human operator. In an example embodiment, the processing device may be configured to cause acceleration of vehicle 5441 via an actuator system that may, for example, include a gas pedal of a vehicle.

[0899] It should be noted that various steps of process 5600 may be modified. Additionally, or alternatively, the sequence of steps of process 5600 may be modified. For example, steps 5607, 5609, and 5611 may precede steps 5603, and 5605.

[0900] Furthermore, although the examples discussed above referred to a crosswalk visible on a road surface, the embodiments may apply to roads in which a crosswalk is not demarcated on the road surface. For example, the disclosed embodiments may identify an intersection and thereby consider one or more areas in a vicinity of the intersection as area(s) (e.g., potential crossing zones) from which pedestrians are likely to travel. Similarly, the disclosed embodiments may determine that a certain location (e.g., based on historical information) is an area (e.g., a potential crossing zone) from which pedestrians are likely to cross a road. This may be based on, for example, observed events in the past (e.g., pedestrians are prone to cross the road at a particular location) and/or based on considerations related to a particular location (e.g., based on building entrances from which pedestrians may emerge or enter).

[0901] The foregoing description has been presented for purposes of illustration. It is not exhaustive and is not limited to the precise forms or embodiments disclosed. Modifications and adaptations will be apparent to those skilled in the art from consideration of the specification and practice of the disclosed embodiments. Additionally, although aspects of the disclosed embodiments are described as being stored in memory, one skilled in the art will appreciate that these aspects can also be stored on other types of computer readable media, such as secondary storage devices, for example, hard disks or CD ROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray, 4K Ultra HD Blu-ray, or other optical drive media.

[0902] Computer programs based on the written description and disclosed methods are within the skill of an experienced developer. The various programs or program modules can be created using any of the techniques known to one skilled in the art or can be designed in connection with existing software. For example, program sections or program modules can be designed in or by means of .Net Framework, .Net Compact Framework (and related languages, such as Visual Basic, C, etc.), Java, C++, Objective-C, HTML, HTML/AJAX combinations, XML, or HTML with included Java applets.

[0903] Moreover, while illustrative embodiments have been described herein, the scope of any and all embodiments having equivalent elements, modifications, omissions, combinations (e.g., of aspects across various embodiments), adaptations and/or alterations as would be appreciated by those skilled in the art based on the present disclosure. The limitations in the claims are to be interpreted broadly based on 5 the language employed in the claims and not limited to examples described in the present specification or during the prosecution of the application. The examples are to be construed as non-exclusive. Furthermore, the steps of the disclosed methods may be modified in any manner, including by reordering steps and/or inserting or deleting steps. It is intended, therefore, that the specification and examples be 10 considered as illustrative only, with a true scope and spirit being indicated by the following claims and their full scope of equivalents.

WHAT IS CLAIMED IS:

1. A system for navigating a host vehicle, the system comprising:
at least one processing device programmed to:

5 receive, from an image capture device, at least one image representative of an
environment of the host vehicle;

determine, based on at least one driving policy, a planned navigational action for
accomplishing a navigational goal of the host vehicle;

10 analyze the at least one image to identify a target vehicle in the environment of
the host vehicle, wherein a direction of travel of the target vehicle is toward the host
vehicle;

determine a next-state distance between the host vehicle and the target vehicle
that would result if the planned navigational action was taken;

15 determine a host vehicle braking rate, a host vehicle maximum acceleration
capability, and a current speed of the host vehicle;

determine a stopping distance for the host vehicle based on the host vehicle
braking rate, the host vehicle maximum acceleration capability, and the current speed of
the host vehicle;

20 determine a current speed of the target vehicle, a target vehicle maximum
acceleration capability, and a target vehicle braking rate;

determine a stopping distance for the target vehicle based on the target vehicle
braking rate, the target vehicle maximum acceleration capability, and the current speed of
the target vehicle; and

25 implement the planned navigational action if the determined next-state distance
is greater than a sum of the stopping distance for the host vehicle and the stopping
distance of the target vehicle.

2. The system of claim 1, wherein the target vehicle and the host vehicle are traveling in a
parking lot.

3. The system of claim 1, wherein the target vehicle travels toward the host vehicle in a lane
occupied by the host vehicle.

30 3. The system of claim 1, wherein the host vehicle braking rate corresponds to a maximum
braking capability of the host vehicle.

4. The system of claim 1, wherein the host vehicle braking rate corresponds to a
predetermined sub-maximal braking capability of the host vehicle.

35 5. The system of claim 1, wherein the target vehicle braking rate corresponds to a maximum
braking capability of the target vehicle determined based on at least one characteristic of the target
vehicle identified through analysis of the at least one image.

6. The system of claim 1, wherein the target vehicle braking rate corresponds to a
predetermined sub-maximal braking capability of the target vehicle.

7. The system of claim 1, wherein the host vehicle braking rate is less than the target vehicle braking rate.

8. The system of claim 1, wherein the current speed of the target vehicle is determined based on analysis of two or more of the received images representative of the environment of the host vehicle.

9. The system of claim 1, wherein the current speed of the target vehicle is determined based on an output of a RADAR unit or a LIDAR unit associated with the host vehicle.

10. The system of claim 1, wherein the stopping distance for the host vehicle includes a host vehicle acceleration distance that corresponds to a distance the host vehicle can travel during a host vehicle reaction time at the host vehicle maximum acceleration capability, starting from the determined current speed of the host vehicle.

11. The system of claim 1, wherein the stopping distance for the target vehicle includes a target vehicle acceleration distance that corresponds to a distance the target vehicle can travel during a target vehicle reaction time at the target vehicle maximum acceleration capability, starting from the determined current speed of the target vehicle.

12. The system of claim 1, wherein at least one of the host vehicle braking rate or the target vehicle braking rate is determined based on a sensed condition of a road surface.

13. The system of claim 1, wherein at least one of the host vehicle braking rate or the target vehicle braking rate is determined based on a sensed weather condition.

14. The system of claim 1, wherein the target vehicle maximum acceleration capability is determined based on at least one characteristic of the target vehicle identified through analysis of the at least one image.

15. The system of claim 14, wherein the at least one characteristic of the target vehicle includes one or more of a vehicle type, a vehicle size, or a vehicle model.

16. The system of claim 1, wherein the planned navigational action includes at least one of a lane change maneuver, a merge maneuver, a passing maneuver, a maintain forward heading maneuver, a parking lot navigation maneuver, or a maintain throttle action.

17. The system of claim 1, wherein the at least one processing device is configured to implement the planned navigational action if the determined next-state distance is greater, by at least a predetermined minimum distance, than the sum of the stopping distance for the host vehicle and the stopping distance of the target vehicle.

18. The system of claim 17, wherein the predetermined minimum distance corresponds to a predetermined separation distance to be maintained between the host vehicle and other vehicles.

19. The system of claim 17, wherein the predetermined separation distance is at least one meter.

20. A method for navigating a host vehicle, the method comprising:
receiving, from an image capture device, at least one image representative of an environment of the host vehicle;

determining, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle;

analyzing the at least one image to identify a target vehicle in the environment of the host vehicle, wherein a direction of travel of the target vehicle is toward the host vehicle;

5 determining a next-state distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken;

determining a host vehicle braking rate, a host vehicle maximum acceleration capability, and a current speed of the host vehicle;

10 determining a stopping distance for the host vehicle based on the host vehicle braking rate, the host vehicle maximum acceleration capability, and the current speed of the host vehicle;

determining a current speed of the target vehicle, a target vehicle maximum acceleration capability, and a target vehicle braking rate;

determining a stopping distance for the target vehicle based on the target vehicle braking rate, the target vehicle maximum acceleration capability, and the current speed of the target vehicle; and

15 implementing the planned navigational action if the determined next-state distance is greater than a sum of the stopping distance for the host vehicle and the stopping distance of the target vehicle.

21. The method of claim 20, wherein the host vehicle braking rate corresponds to a maximum braking capability of the host vehicle.

20 22. The method of claim 20, wherein the host vehicle braking rate corresponds to a predetermined sub-maximal braking capability of the host vehicle.

23. The method of claim 20, wherein the target vehicle braking rate corresponds to a maximum braking capability of the target vehicle determined based on at least one characteristic of the target vehicle identified through analysis of the at least one image.

25 24. The method of claim 20, wherein the target vehicle braking rate corresponds to a predetermined sub-maximal braking capability of the target vehicle.

25. The method of claim 20, wherein the host vehicle braking rate is less than the target vehicle braking rate.

30 26. The method of claim 20, wherein the current speed of the target vehicle is determined based on analysis of two or more of the received images representative of the environment of the host vehicle.

27. The method of claim 20, wherein the current speed of the target vehicle is determined based on an output of a RADAR unit or a LIDAR unit associated with the host vehicle.

35 28. The method of claim 20, wherein the target vehicle maximum acceleration capability is determined based on at least one characteristic of the target vehicle identified through analysis of the at least one image.

29. The method of claim 28, wherein the at least one characteristic of the target vehicle includes one or more of a vehicle type, a vehicle size, or a vehicle model.

30. The method of claim 20, wherein the planned navigational action includes at least one of a lane change maneuver, a merge maneuver, a passing maneuver, a maintain forward heading maneuver, a parking lot navigation maneuver, or a maintain throttle action.

31. A system for navigating a host vehicle, the system comprising:

5 at least one processing device programmed to:

receive, from an image capture device, at least one image representative of an environment of the host vehicle;

determine, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle;

10 analyze the at least one image to identify a target vehicle in the environment of the host vehicle;

determine a next-state lateral distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken;

15 determine a maximum yaw rate capability of the host vehicle, a maximum change in turn radius capability of the host vehicle, and a current lateral speed of the host vehicle;

determine a lateral braking distance for the host vehicle based on the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, and the current lateral speed of the host vehicle;

20 determine a current lateral speed of the target vehicle, a target vehicle maximum yaw rate capability, and a target vehicle maximum change in turn radius capability;

determine a lateral braking distance for the target vehicle based on the current lateral speed of the target vehicle, the target vehicle maximum yaw rate capability, and the target vehicle maximum change in turn radius capability; and

25 implement the planned navigational action if the determined next-state lateral distance is greater than a sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle.

32. The system of claim 31, wherein the lateral braking distance for the host vehicle corresponds to a distance required to achieve a host vehicle lateral speed of zero, at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle, starting from the current lateral speed of the host vehicle and after a lateral acceleration of the host vehicle as a result of a turning maneuver made toward the target vehicle at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle during a reaction time associated with the host vehicle.

35 33. The system of claim 31, wherein the lateral braking distance for the target vehicle corresponds to a distance required to achieve a target vehicle lateral speed of zero, at the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability, starting from the current lateral speed of the target vehicle and after a lateral acceleration of the target vehicle as a

result of a turning maneuver by the target vehicle toward the host vehicle made at the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability during a reaction time associated with the target vehicle.

34. The system of claim 31, wherein the at least one processing device is programmed to
5 implement the planned navigational action regardless of the next-state lateral distance if the planned navigational action would result in a safe longitudinal distance between the host vehicle and the target vehicle.

35. The system of claim 34, wherein the safe longitudinal distance is greater than or equal to
10 a sum of a host vehicle longitudinal stopping distance and a target vehicle longitudinal stopping distance;
wherein the host vehicle stopping distance is based on a maximum longitudinal braking capability
of the host vehicle and a current longitudinal speed of the host vehicle, and wherein the host vehicle
stopping distance includes a host vehicle acceleration distance corresponding to a distance the host
vehicle can travel at a maximum acceleration capability of the host vehicle during a reaction time
associated with the host vehicle starting from the current longitudinal speed of the host vehicle; and
15 wherein the target vehicle stopping distance is based on a maximum longitudinal braking
capability of the target vehicle and a current longitudinal speed of the target vehicle, and wherein the
target vehicle stopping distance includes a target vehicle acceleration distance corresponding to a distance
the target vehicle can travel at a maximum acceleration capability of the target vehicle during a reaction
time associated with the target vehicle starting from the current longitudinal speed of the target vehicle.

36. The system of claim 31, wherein at least one of the target vehicle maximum yaw rate
20 capability and the target vehicle maximum change in turn radius capability is determined based on at least
one characteristic of the target vehicle identified through analysis of the at least one image.

37. The system of claim 6, wherein the at least one characteristic of the target vehicle
includes one or more of a vehicle type, a vehicle size, or a vehicle model.

38. The system of claim 31, wherein the maximum yaw rate capability of the host vehicle
25 and the maximum change in turn radius capability of the host vehicle are based on predetermined
constraints for the host vehicle.

39. The system of claim 31, wherein the current lateral speed of the target vehicle is
determined based on an output of a RADAR unit or a LIDAR unit associated with the host vehicle.

40. The system of claim 31, wherein at least one of the maximum yaw rate capability of the
30 host vehicle or the maximum change in turn radius capability of the host vehicle is determined based on a
sensed condition of a road surface.

41. The system of claim 31, wherein at least one of the target vehicle maximum yaw rate
35 capability or the target vehicle maximum change in turn radius capability is determined based on a sensed
condition of a road surface.

42. The system of claim 31, wherein at least one of the maximum yaw rate capability of the
host vehicle, the maximum change in turn radius capability of the host vehicle, the target vehicle

maximum yaw rate capability, or the target vehicle maximum change in turn radius capability is determined based on a sensed weather condition.

43. The system of claim 31, wherein the planned navigational action includes at least one of a lane change maneuver, a merge maneuver, a passing maneuver, a parking lot navigation maneuver, or a turning maneuver.

44. The system of claim 31, wherein the at least one processing device is configured to implement the planned navigational action if the determined next-state lateral distance is greater, by at least a predetermined minimum distance, than the sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle.

45. The system of claim 44, wherein the predetermined minimum distance corresponds to a predetermined lateral separation distance to be maintained between the host vehicle and other vehicles.

46. The system of claim 15, wherein the predetermined lateral separation distance is at least one meter.

47. A method for navigating a host vehicle, the method comprising:

receiving, from an image capture device, at least one image representative of an environment of the host vehicle;

determining, based on at least one driving policy, a planned navigational action for accomplishing a navigational goal of the host vehicle;

analyzing the at least one image to identify a target vehicle in the environment of the host vehicle;

determining a next-state lateral distance between the host vehicle and the target vehicle that would result if the planned navigational action was taken;

determining a maximum yaw rate capability of the host vehicle, a maximum change in turn radius capability of the host vehicle, and a current lateral speed of the host vehicle;

determining a lateral braking distance for the host vehicle based on the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, and the current lateral speed of the host vehicle;

determining a current lateral speed of the target vehicle, a target vehicle maximum yaw rate capability, and a target vehicle maximum change in turn radius capability;

determining a lateral braking distance for the target vehicle based on the current lateral speed of the target vehicle, the target vehicle maximum yaw rate capability, and the target vehicle maximum change in turn radius capability; and

implementing the planned navigational action if the determined next-state lateral distance is greater than a sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle.

48. The method of claim 47, wherein the lateral braking distance for the host vehicle corresponds to a distance required to achieve a host vehicle lateral speed of zero, at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle, starting from the current lateral speed of the host vehicle and after a lateral acceleration of the host

vehicle as a result of a turning maneuver made toward the target vehicle at the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle during a reaction time associated with the host vehicle.

49. The method of claim 47, wherein the lateral braking distance for the target vehicle
5 corresponds to a distance required to achieve a target vehicle lateral speed of zero, at the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability, starting from the current lateral speed of the target vehicle and after a lateral acceleration of the target vehicle as a result of a turning maneuver by the target vehicle toward the host vehicle made at the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability during a
10 reaction time associated with the target vehicle.

50. The method of claim 47, wherein the at least one processing device is programmed to implement the planned navigational action regardless of the next-state lateral distance if the planned navigational action would result in a safe longitudinal distance between the host vehicle and the target vehicle.

51. The method of claim 50, wherein the safe longitudinal distance is greater than or equal to
15 a sum of a host vehicle longitudinal stopping distance and a target vehicle longitudinal stopping distance; wherein the host vehicle stopping distance is based on a maximum longitudinal braking capability of the host vehicle and a current longitudinal speed of the host vehicle, and wherein the host vehicle stopping distance includes a host vehicle acceleration distance corresponding to a distance the host
20 vehicle can travel at a maximum acceleration capability of the host vehicle during a reaction time associated with the host vehicle starting from the current longitudinal speed of the host vehicle; and wherein the target vehicle stopping distance is based on a maximum longitudinal braking capability of the target vehicle and a current longitudinal speed of the target vehicle, and wherein the target vehicle stopping distance includes a target vehicle acceleration distance corresponding to a distance
25 the target vehicle can travel at a maximum acceleration capability of the target vehicle during a reaction time associated with the target vehicle starting from the current longitudinal speed of the target vehicle.

52. The method of claim 47, wherein at least one of the target vehicle maximum yaw rate capability and the target vehicle maximum change in turn radius capability is determined based on at least one characteristic of the target vehicle identified through analysis of the at least one image.

53. The system of claim 52, wherein the at least one characteristic of the target vehicle
30 includes one or more of a vehicle type, a vehicle size, or a vehicle model.

54. The method of claim 47, wherein the maximum yaw rate capability of the host vehicle and the maximum change in turn radius capability of the host vehicle are based on predetermined constraints for the host vehicle.

55. The method of claim 47, wherein the current lateral speed of the target vehicle is
35 determined based on an output of a RADAR unit or a LIDAR unit associated with the host vehicle.

56. The method of claim 47, wherein at least one of the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, the target vehicle

maximum yaw rate capability, or the target vehicle maximum change in turn radius capability is determined based on a sensed condition of a road surface.

57. The method of claim 47, wherein at least one of the maximum yaw rate capability of the host vehicle, the maximum change in turn radius capability of the host vehicle, the target vehicle maximum yaw rate capability, or the target vehicle maximum change in turn radius capability is determined based on a sensed weather condition.

58. The method of claim 47, wherein the planned navigational action includes at least one of a lane change maneuver, a merge maneuver, a passing maneuver, a parking lot navigation maneuver, or a turning maneuver.

59. The method of claim 47, wherein the at least one processing device is configured to implement the planned navigational action if the determined next-state lateral distance is greater, by at least a predetermined minimum distance, than the sum of the lateral braking distance for the host vehicle and the lateral braking distance of the target vehicle.

60. The method of claim 59, wherein the predetermined minimum distance corresponds to a predetermined lateral separation distance to be maintained between the host vehicle and other vehicles.

61. A system for navigating a host vehicle in proximity to a pedestrian crosswalk, the system comprising at least one processing device programmed to:
receive, from an image capture device, at least one image representative of an environment of the host vehicle;
detect, based on analysis of the at least one image, a representation of a pedestrian crosswalk in the at least one image;
determine, based on analysis of the at least one image, whether a representation of a pedestrian appears in the at least one image;
detect a presence of a traffic light in the environment of the host vehicle;
determine whether the detected traffic light is relevant to the host vehicle and the pedestrian crosswalk;
determine a state of the detected traffic light;
determine, when a representation of a pedestrian appears in the at least one image, a proximity of the pedestrian relative to the detected pedestrian crosswalk;
determine, based on at least one driving policy, a planned navigational action for causing the host vehicle to navigate relative to the detected pedestrian crosswalk, wherein determination of the planned navigational action is further based on the determined state of the detected traffic light and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk; and
cause one or more actuator systems of the host vehicle to implement the planned navigational action.

62. The system of claim 61, wherein the planned navigational action for the host vehicle includes traversing the pedestrian crosswalk without a pedestrian-induced adjustment when the detected

traffic light is in a green state and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk indicates that the pedestrian is not likely to enter the crosswalk prior to an intersection of the host vehicle with the crosswalk.

- 5 63. The system of claim 61, wherein the planned navigational action for the host vehicle includes traversing the pedestrian crosswalk without a pedestrian-induced adjustment when the detected traffic light is in a green state and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk, in combination with a detected heading and speed associated with the pedestrian, indicates that the pedestrian cannot enter the crosswalk prior to an intersection of the host vehicle with the crosswalk.
- 10 64. The system of claim 61, wherein the planned navigational action for the host vehicle includes traversing the pedestrian crosswalk with a pedestrian-induced adjustment when the detected traffic light is in a green state and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk indicates that the pedestrian may enter the crosswalk prior to an intersection of the host vehicle with the crosswalk.
- 15 65. The system of claim 64, wherein the pedestrian-induced adjustment is determined based on a detected heading and speed of the pedestrian.
66. The system of claim 64, wherein the pedestrian-induced adjustment includes a lateral movement away from the pedestrian.
67. The system of claim 64, wherein the pedestrian-induced adjustment includes a slowing of the host vehicle prior to intersection of the host vehicle with the pedestrian crosswalk.
- 20 68. The system of claim 61, wherein the planned navigational action for the host vehicle includes stopping the host vehicle prior to the pedestrian crosswalk, when the traffic light is in a green state and the pedestrian is detected in the pedestrian crosswalk.
69. The system of claim 61, wherein the determination of whether the detected traffic light is relevant to the host vehicle and the pedestrian crosswalk is based on traffic light relevancy information included in one or maps accessible by the at least one processing device.
- 25 70. The system of claim 61, wherein detection of the presence of the traffic light in the environment of the host vehicle is based on analysis of the at least one image.
71. The system of claim 61, wherein detection of the presence of the traffic light in the environment of the host vehicle is based on one or more signals transmitted by the detected traffic light.
- 30 72. The system of claim 61, wherein determination of the state of the detected traffic light is based on analysis of the at least one image.
73. The system of claim 61, wherein determination of the state of the detected traffic light is based on one or more signals transmitted by the detected traffic light.
- 35 74. A method for navigating a host vehicle in proximity to a pedestrian crosswalk, the method comprising:
receiving, from an image capture device, at least one image representative of an environment of the host vehicle;

detecting, based on analysis of the at least one image, a representation of a pedestrian crosswalk in the at least one image;
determining, based on analysis of the at least one image, whether a representation of a pedestrian appears in the at least one image;
5 detecting a presence of a traffic light in the environment of the host vehicle;
determining whether the detected traffic light is relevant to the host vehicle and the pedestrian crosswalk;
determining a state of the detected traffic light;
determining, when a representation of a pedestrian appears in the at least one image, a proximity
10 of the pedestrian relative to the detected pedestrian crosswalk;
determining, based on at least one driving policy, a planned navigational action for causing the host vehicle to navigate relative to the detected pedestrian crosswalk, wherein determination of the planned navigational action is further based on the determined state of the detected traffic light and the determined proximity of the pedestrian relative to the
15 detected pedestrian crosswalk; and
causing one or more actuator systems of the host vehicle to implement the planned navigational action.

75. The method of claim 74, wherein detection of the presence of the traffic light in the environment of the host vehicle is based on one or more signals transmitted by the detected traffic light.

20 76. The method of claim 74, wherein the planned navigational action for the host vehicle includes traversing the pedestrian crosswalk without a pedestrian-induced adjustment when the detected traffic light is in a green state and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk indicates that the pedestrian is not likely to enter the crosswalk prior to an intersection of the host vehicle with the crosswalk.

25 77. The method of claim 74, wherein the planned navigational action for the host vehicle includes traversing the pedestrian crosswalk without a pedestrian-induced adjustment when the detected traffic light is in a green state and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk, in combination with a detected heading and speed associated with the pedestrian, indicates that the pedestrian cannot enter the crosswalk prior to an intersection of
30 the host vehicle with the crosswalk.

78. The method of claim 74, wherein the planned navigational action for the host vehicle includes traversing the pedestrian crosswalk with a pedestrian-induced adjustment when the detected traffic light is in a green state and the determined proximity of the pedestrian relative to the detected pedestrian crosswalk indicates that the pedestrian may enter the crosswalk prior to an
35 intersection of the host vehicle with the crosswalk.

79. A method for navigating a vehicle in proximity to a crosswalk, the method comprising:
receiving, from an image capture device, at least one image representing an environment of the
host vehicle;

detecting a starting and an ending location of a crosswalk;
determining, based on an analysis of the at least one image, whether a pedestrian is present in a
proximity of the crosswalk;

detecting a presence of a traffic light in the environment of the host vehicle;

5 determining whether the traffic light is relevant to the host vehicle and the crosswalk;

determining a state of the traffic light;

determining a navigational action for the host vehicle in vicinity of the detected crosswalk, based
on:

the relevancy of the traffic light;

10 the determined state of the traffic light;

a presence of the pedestrian in the proximity of the crosswalk;

a shortest distance selected between the pedestrian and one of the starting or the ending
location of the crosswalk; and

a motion vector of the pedestrian; and

15 causing one or more actuator systems of the host vehicle to implement the navigational action.

80. The method of claim 79, wherein detection of the presence of the traffic light in the environment
of the host vehicle is based on one or more signals transmitted by the detected traffic light.




FIG. 1




FIG. 2A




FIG. 2B




FIG. 2C




FIG. 2D




FIG. 2E

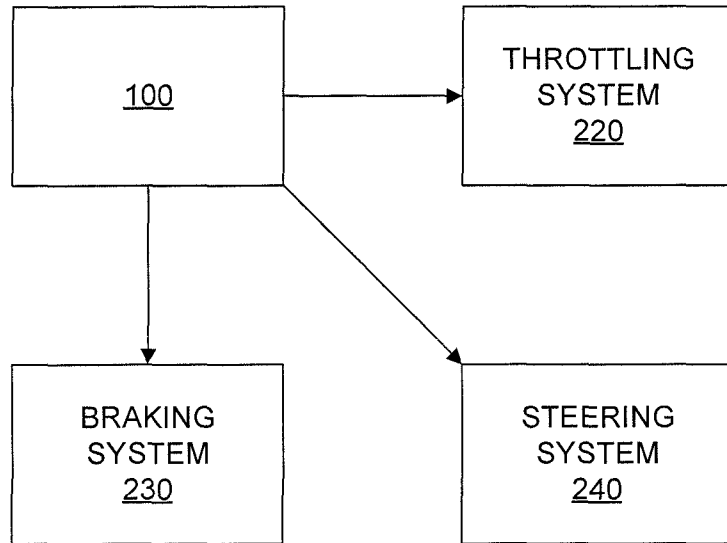


FIG. 2F




FIG. 3A




FIG. 3B




FIG. 3C




FIG. 3D


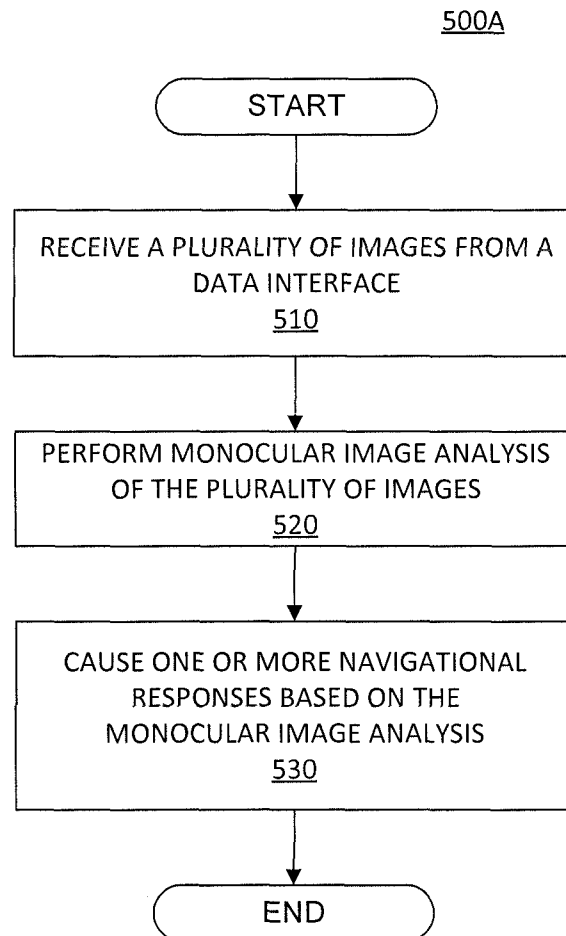


FIG. 4

**FIG. 5A**

500B

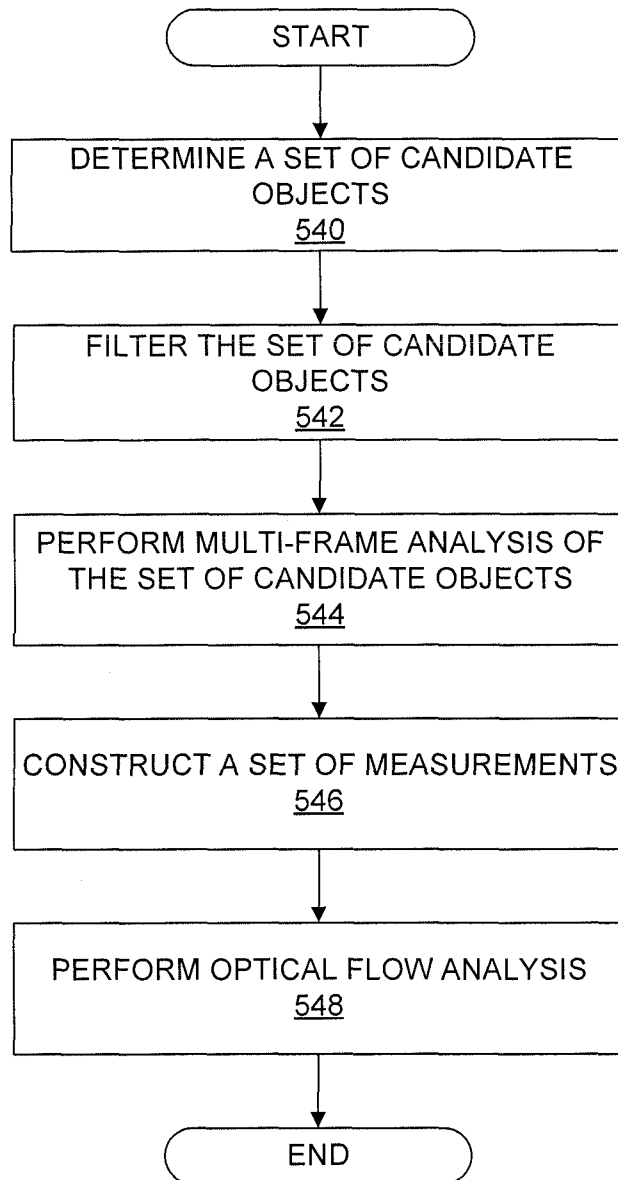
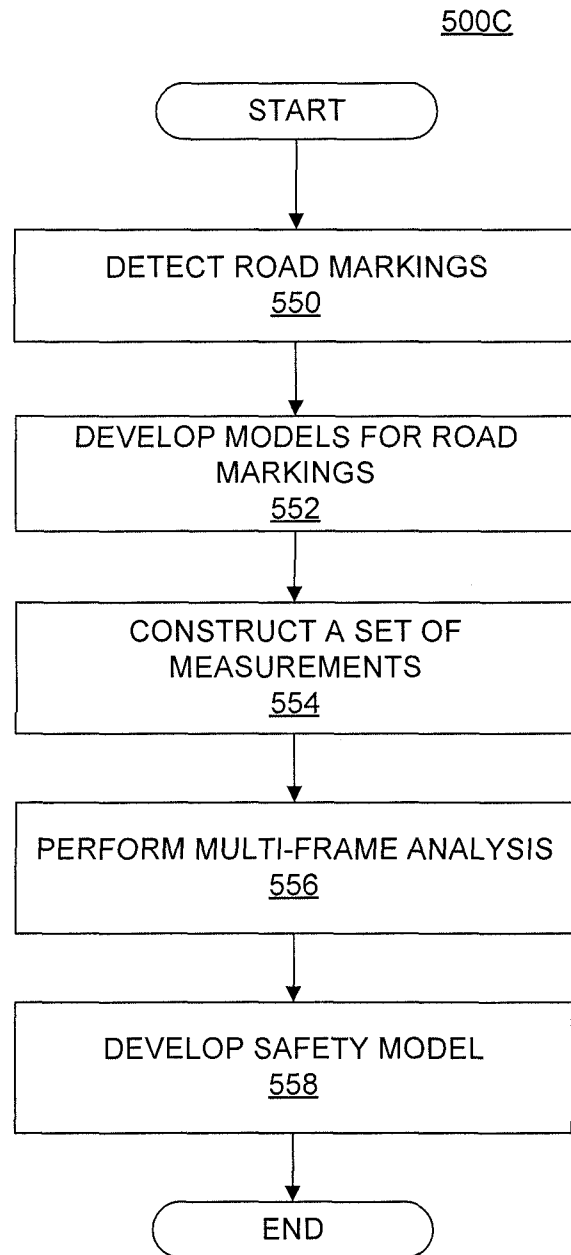


FIG. 5B

**FIG. 5C**

500D

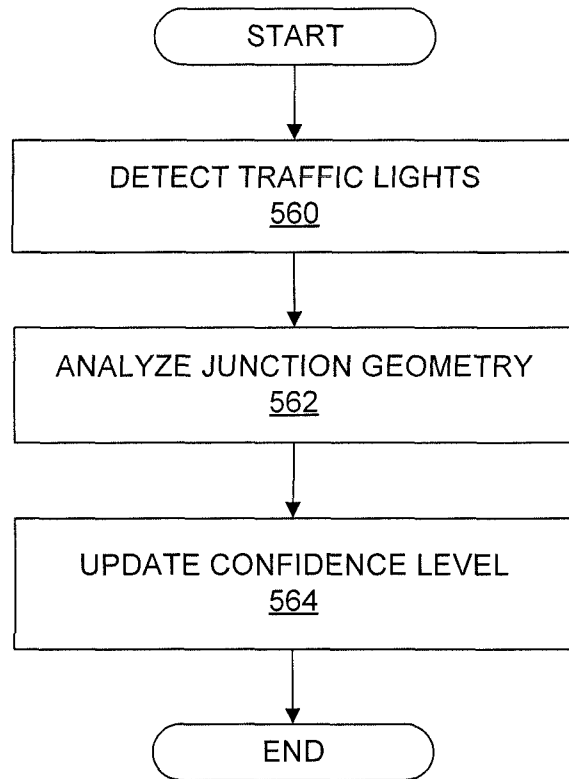


FIG. 5D

500E

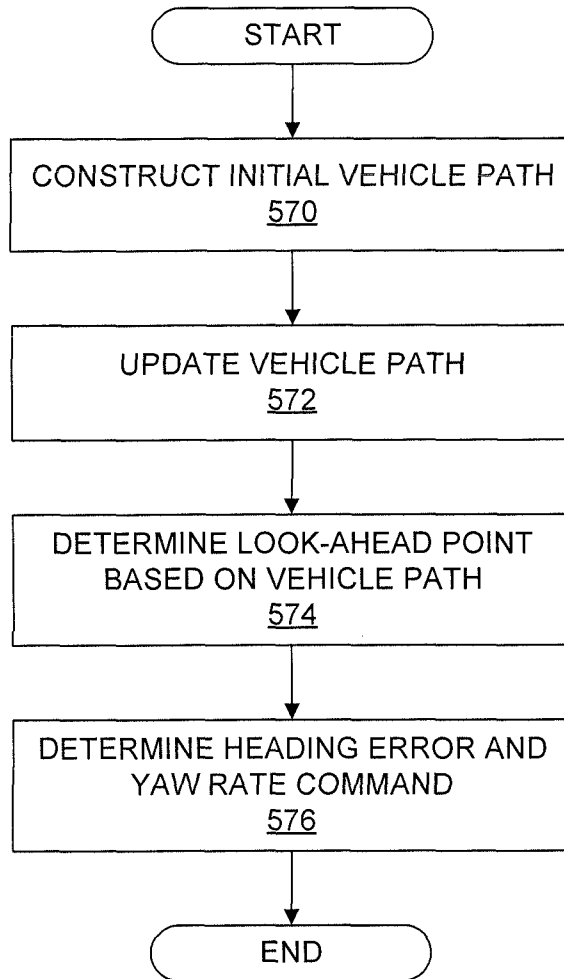


FIG. 5E

500F

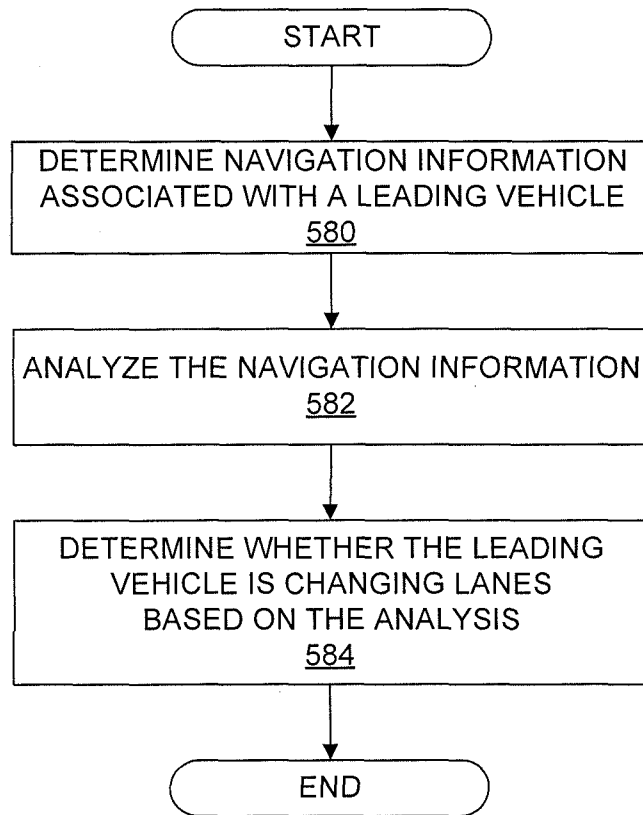
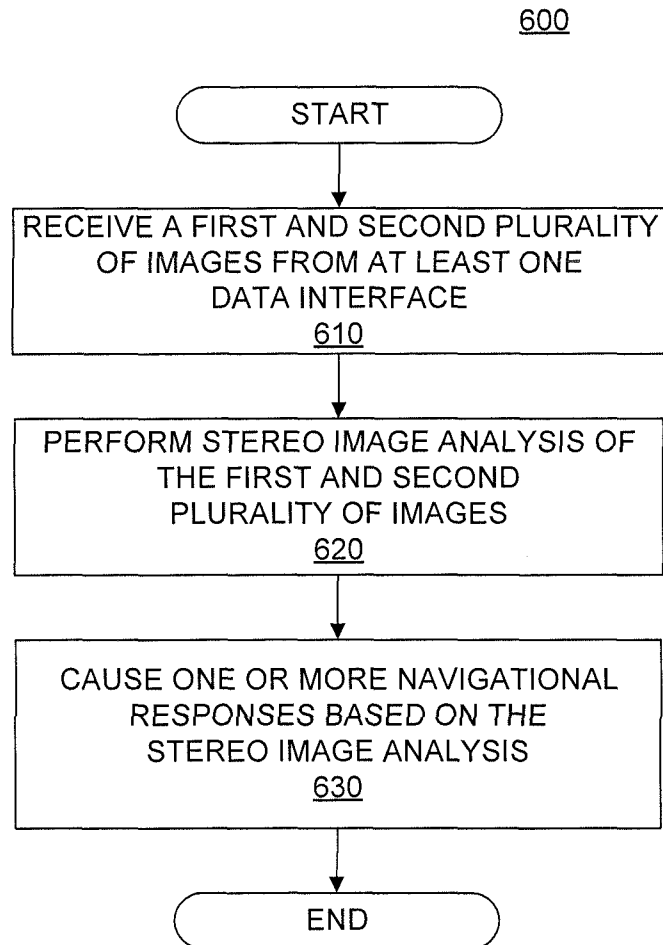
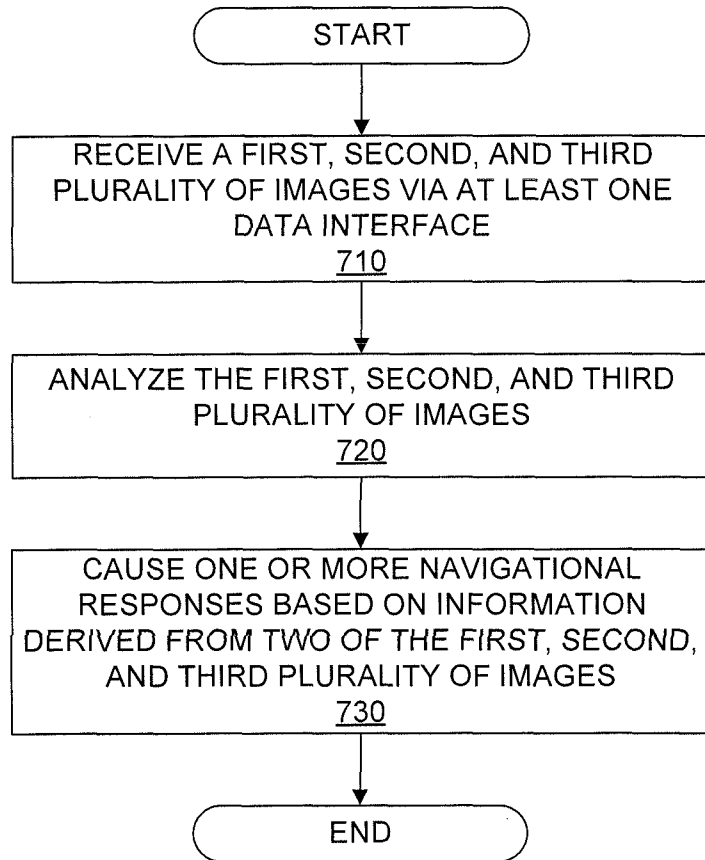


FIG. 5F

**FIG. 6**

700**FIG. 7**




FIG. 8

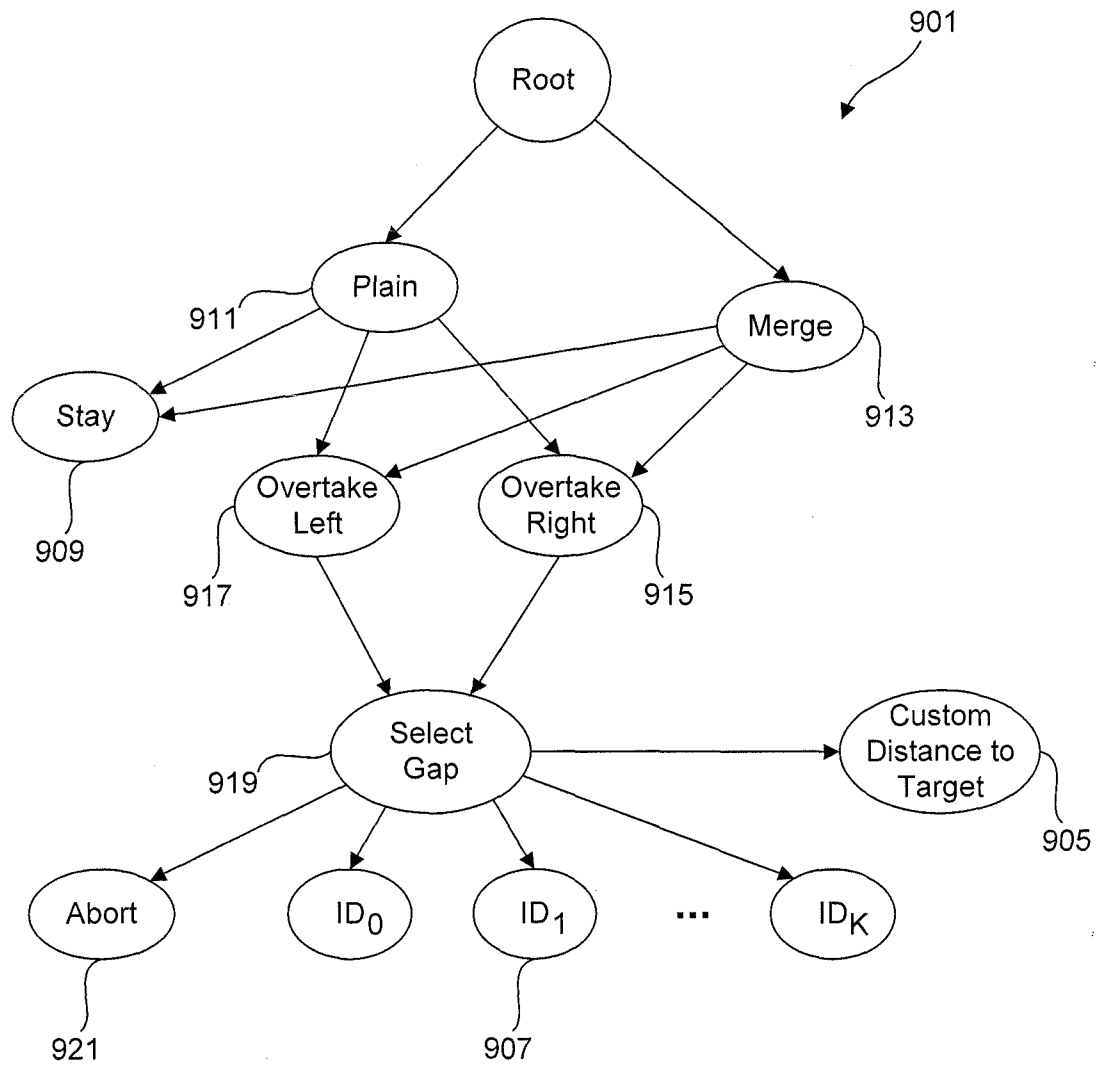


FIG. 9

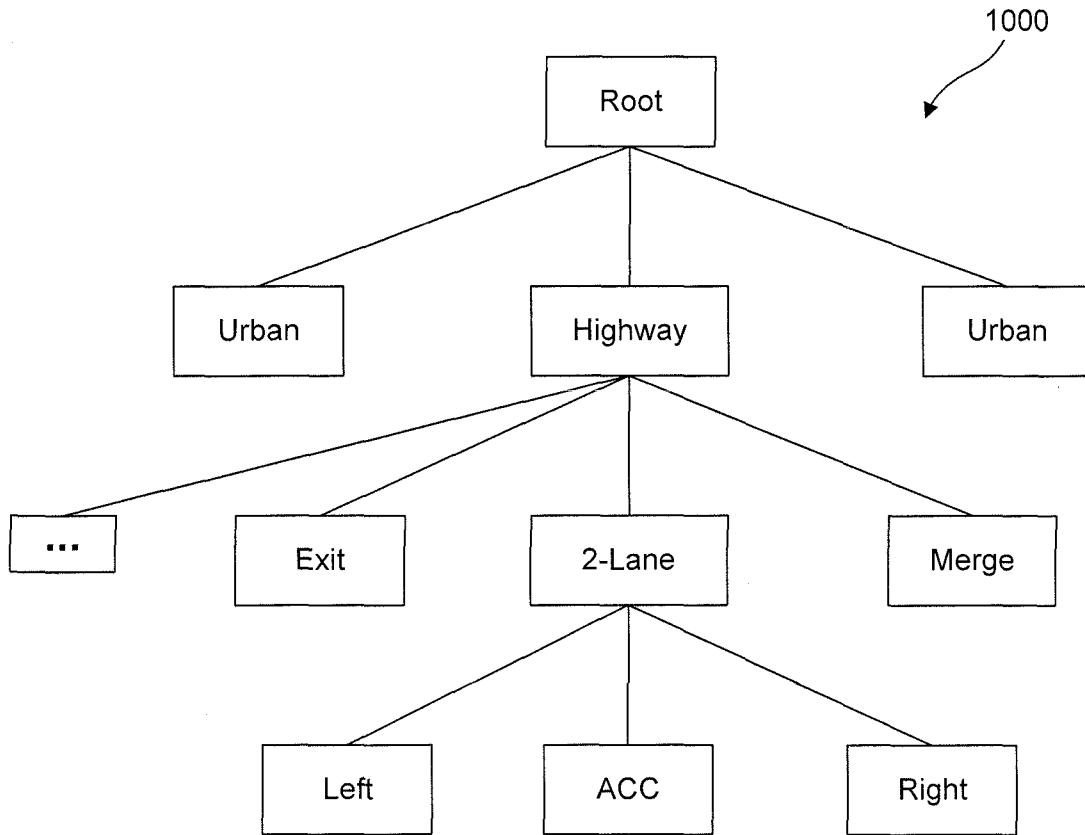


FIG. 10




FIG. 11A




FIG. 11B




FIG. 11C




FIG. 11D




FIG. 11E




FIG. 12

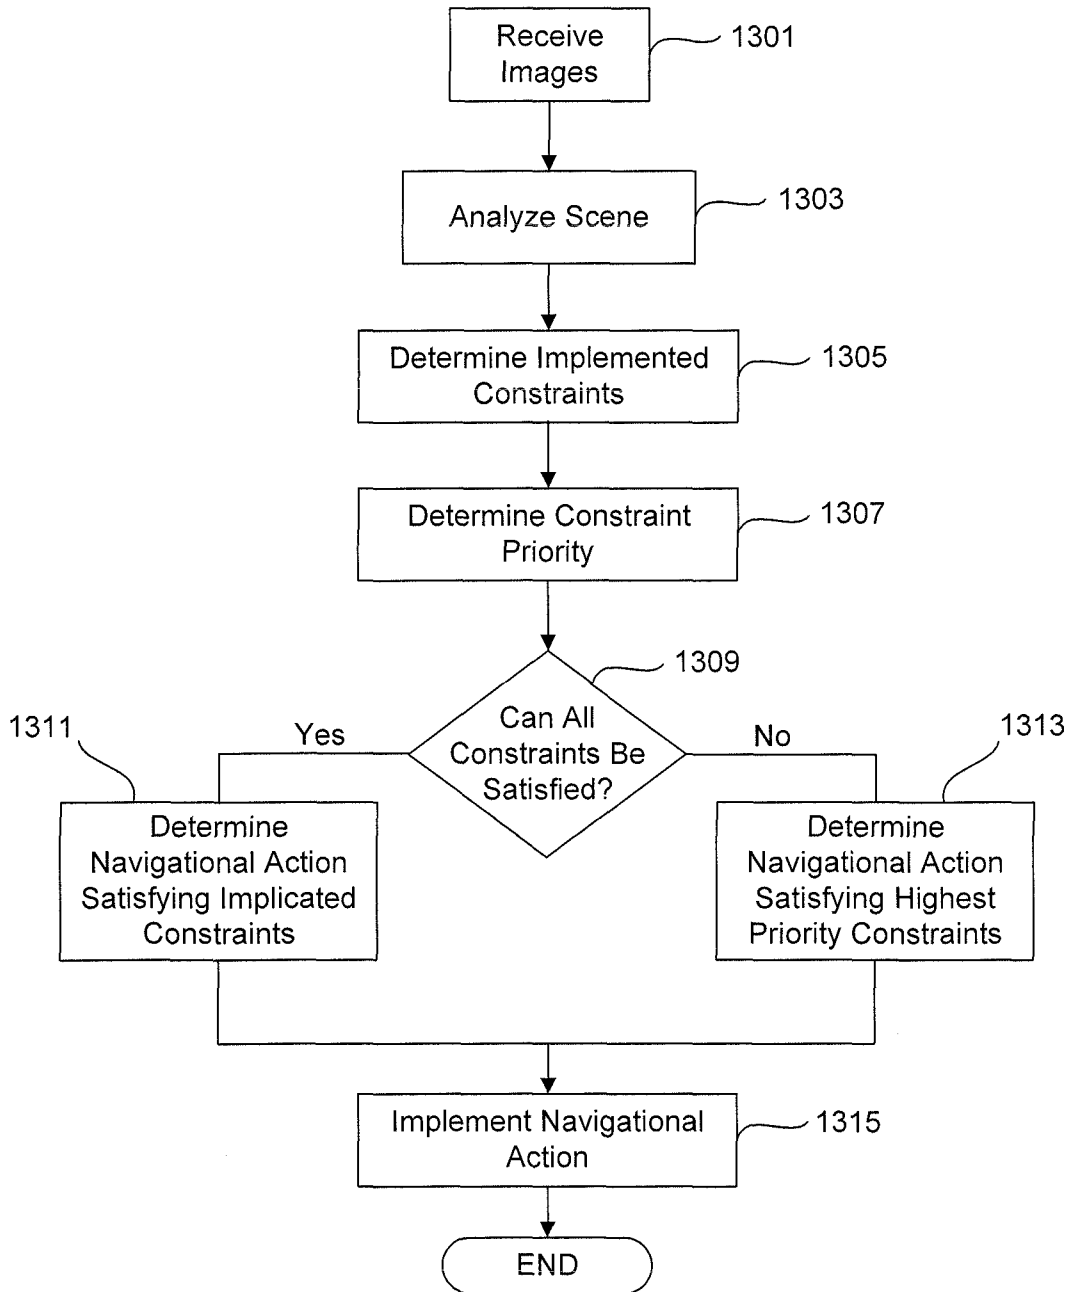
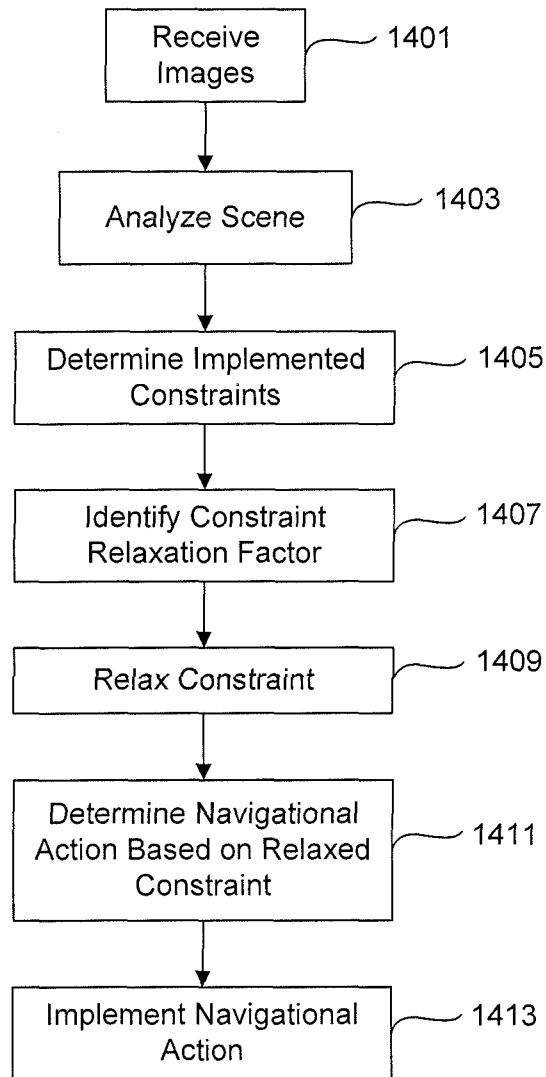
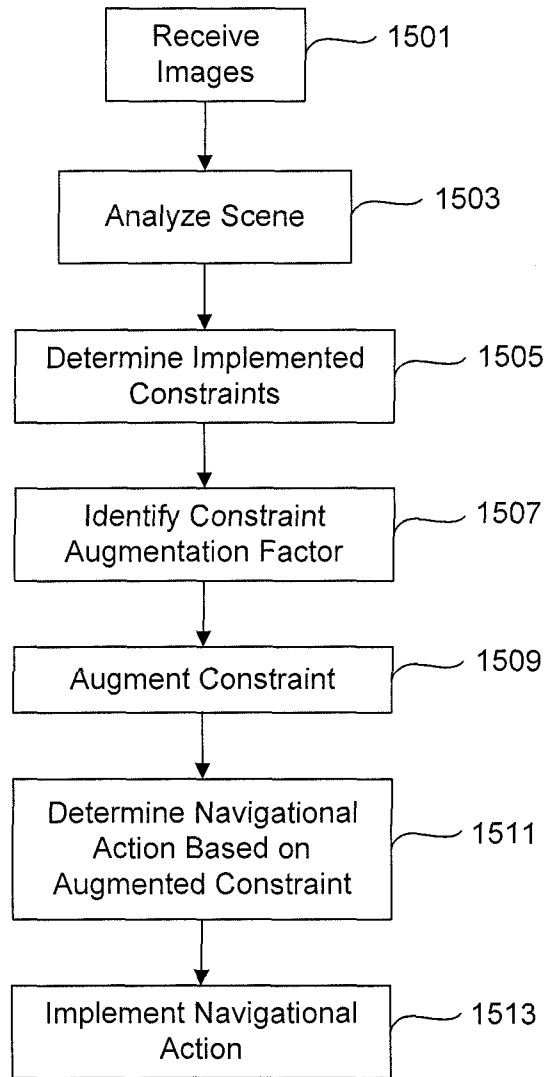


FIG. 13

**FIG. 14**

**FIG. 15**

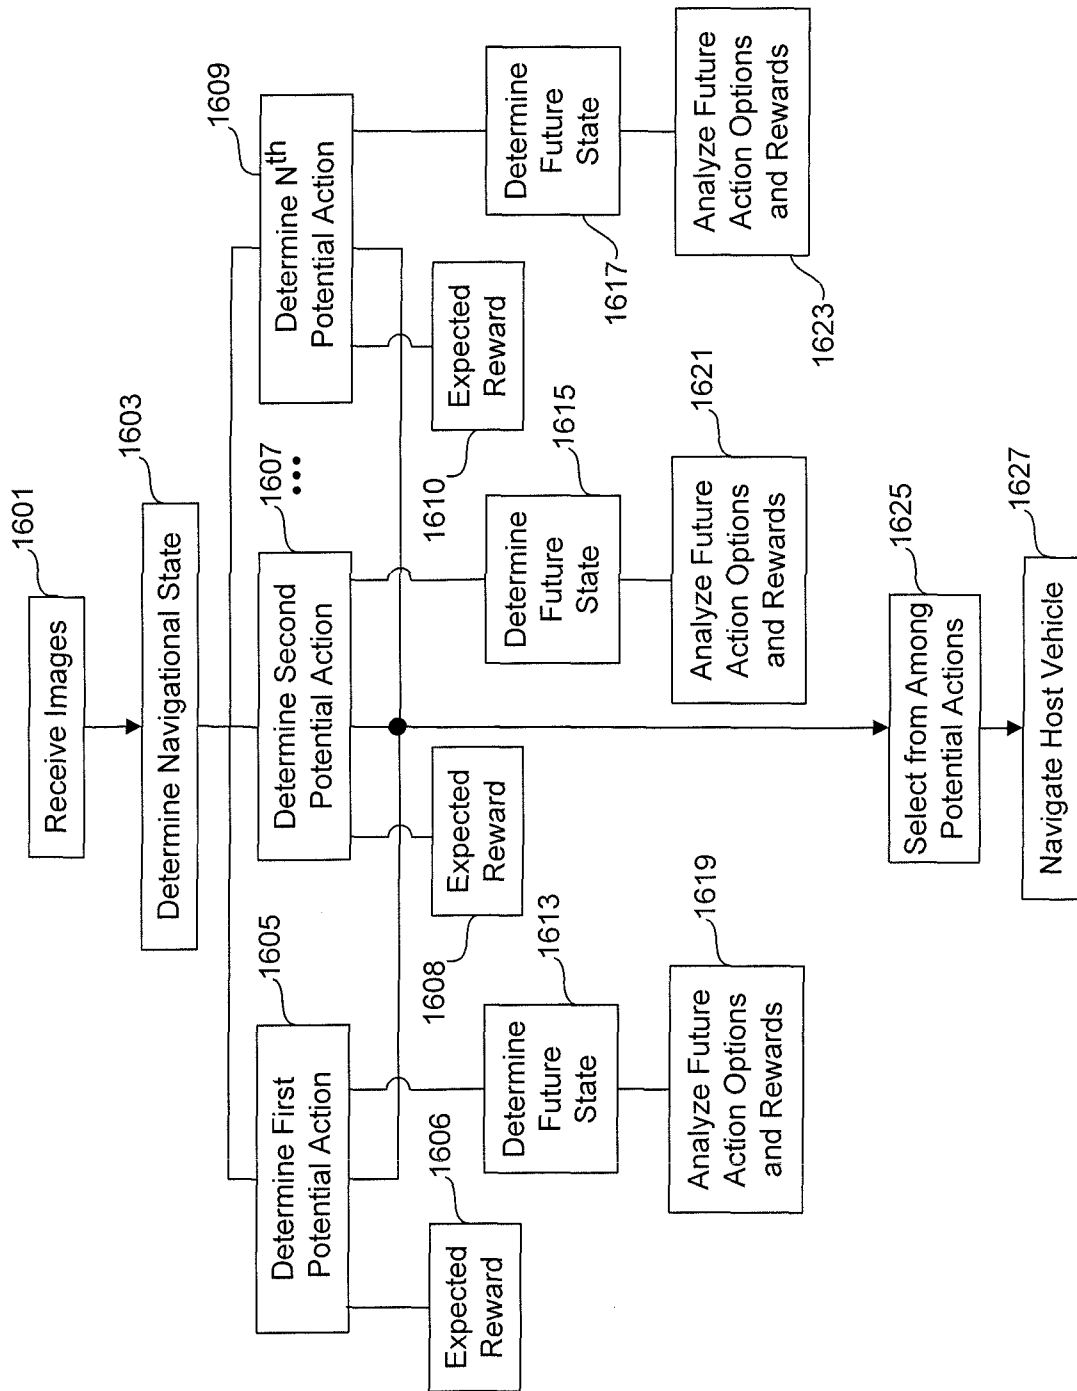


FIG. 16




FIG. 17A




FIG. 17B

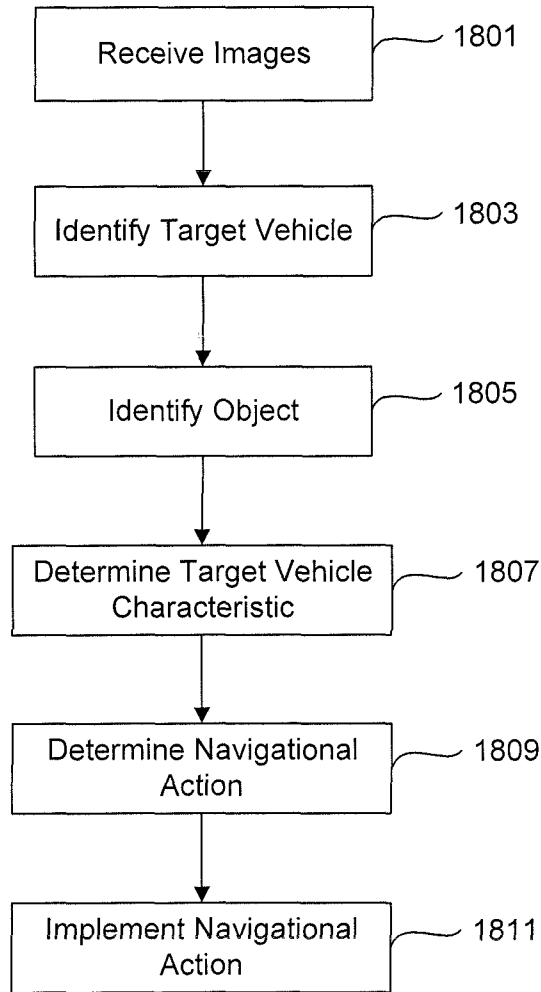


FIG. 18




Fig. 19




Fig. 20A




Fig. 20B




Fig. 21




Fig. 22




Fig. 23




Fig. 24D

Fig. 24C

Fig. 24B

Fig. 24A




Fig. 25




Fig. 26




Fig. 27




Fig. 28A




Fig. 28B




Fig. 29A




Fig. 29B




Fig. 30A




Fig. 30B




Fig. 31A




Fig. 31B




Fig. 31C




Fig. 31D




Fig. 32A




Fig. 32B




Fig. 33A




Fig. 33B




Fig. 34A




Fig. 34B




Fig. 35A




Fig. 35B




Fig. 36A




Fig. 36B




Fig. 37A




Fig. 37B




Fig. 38A




Fig. 38B




Fig. 39A




Fig. 39B




Fig. 40A




Fig. 40B




Fig. 41A




Fig. 41B




Fig. 42A




Fig. 42B




Fig. 43A




Fig. 43B




Fig. 43C




Fig. 44A




Fig. 44B




Fig. 44C




Fig. 45A




Fig. 45B




Fig. 45C




Fig. 46A




Fig. 46B




Fig. 46C




Fig. 46D




FIG. 47A




FIG. 47B




FIG. 47C




FIG. 47D




FIG. 47E




FIG. 47F





FIG. 47G




FIG. 47H




FIG. 471




FIG. 48A




FIG. 48B




FIG. 49

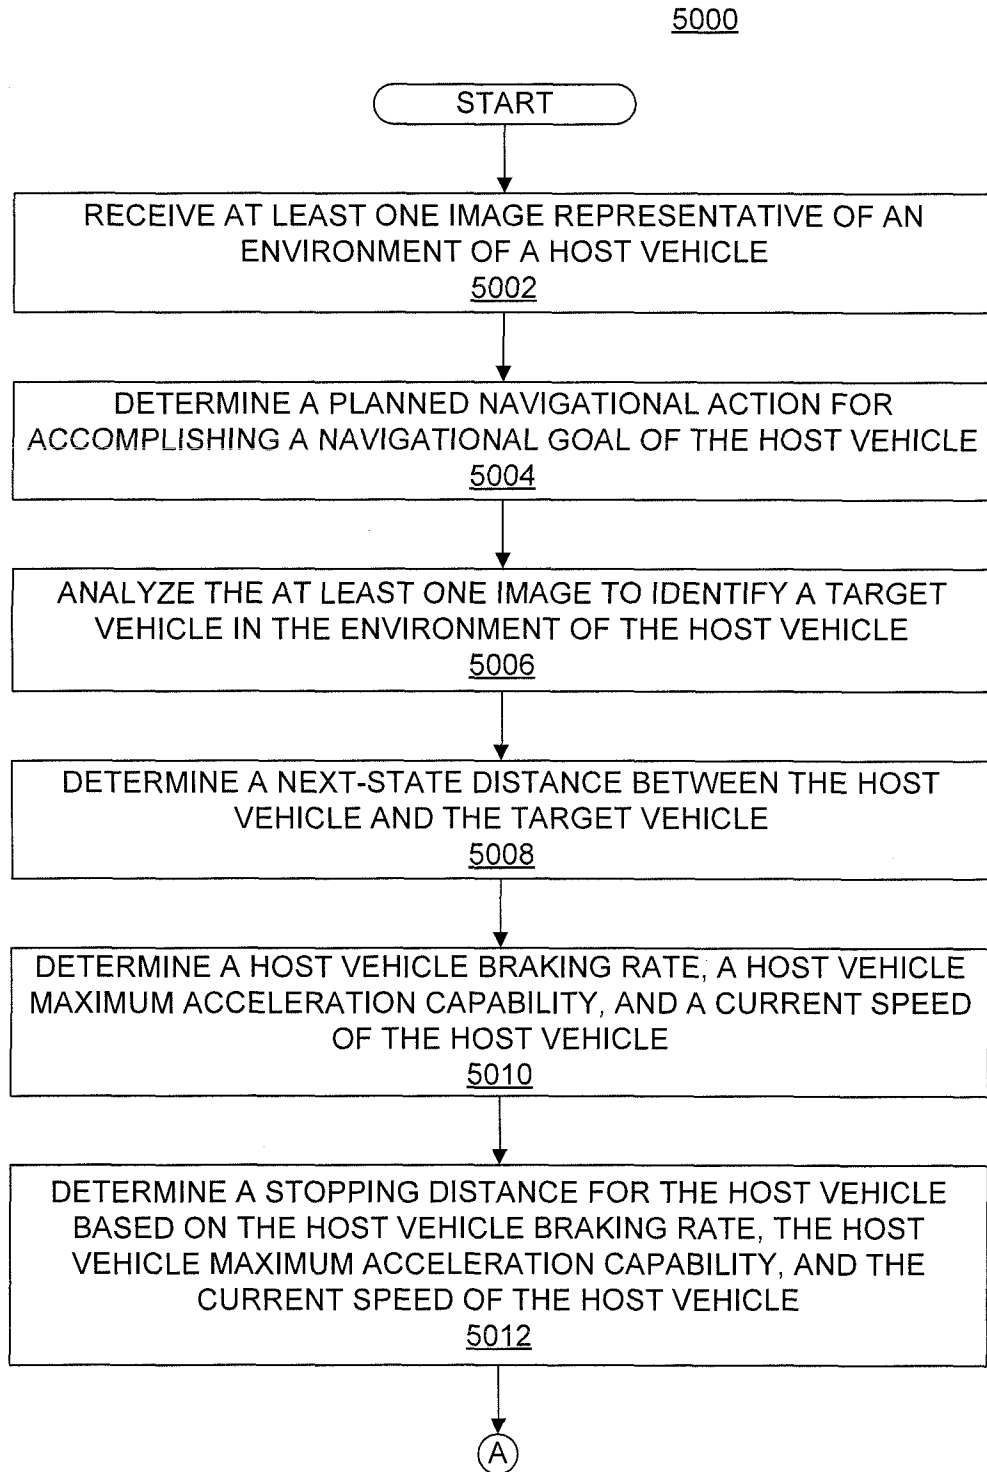


FIG. 50A

5000

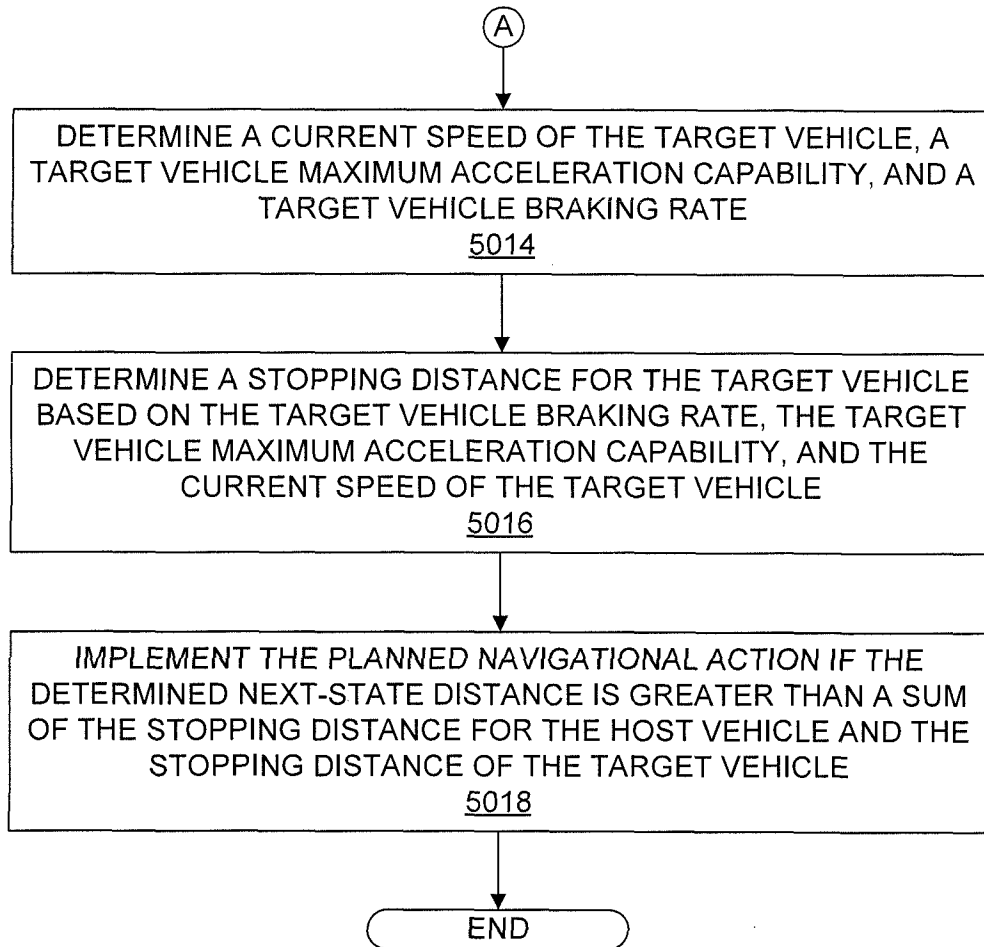


FIG. 50B




FIG. 51A




FIG. 51B




FIG. 52A




FIG. 52B

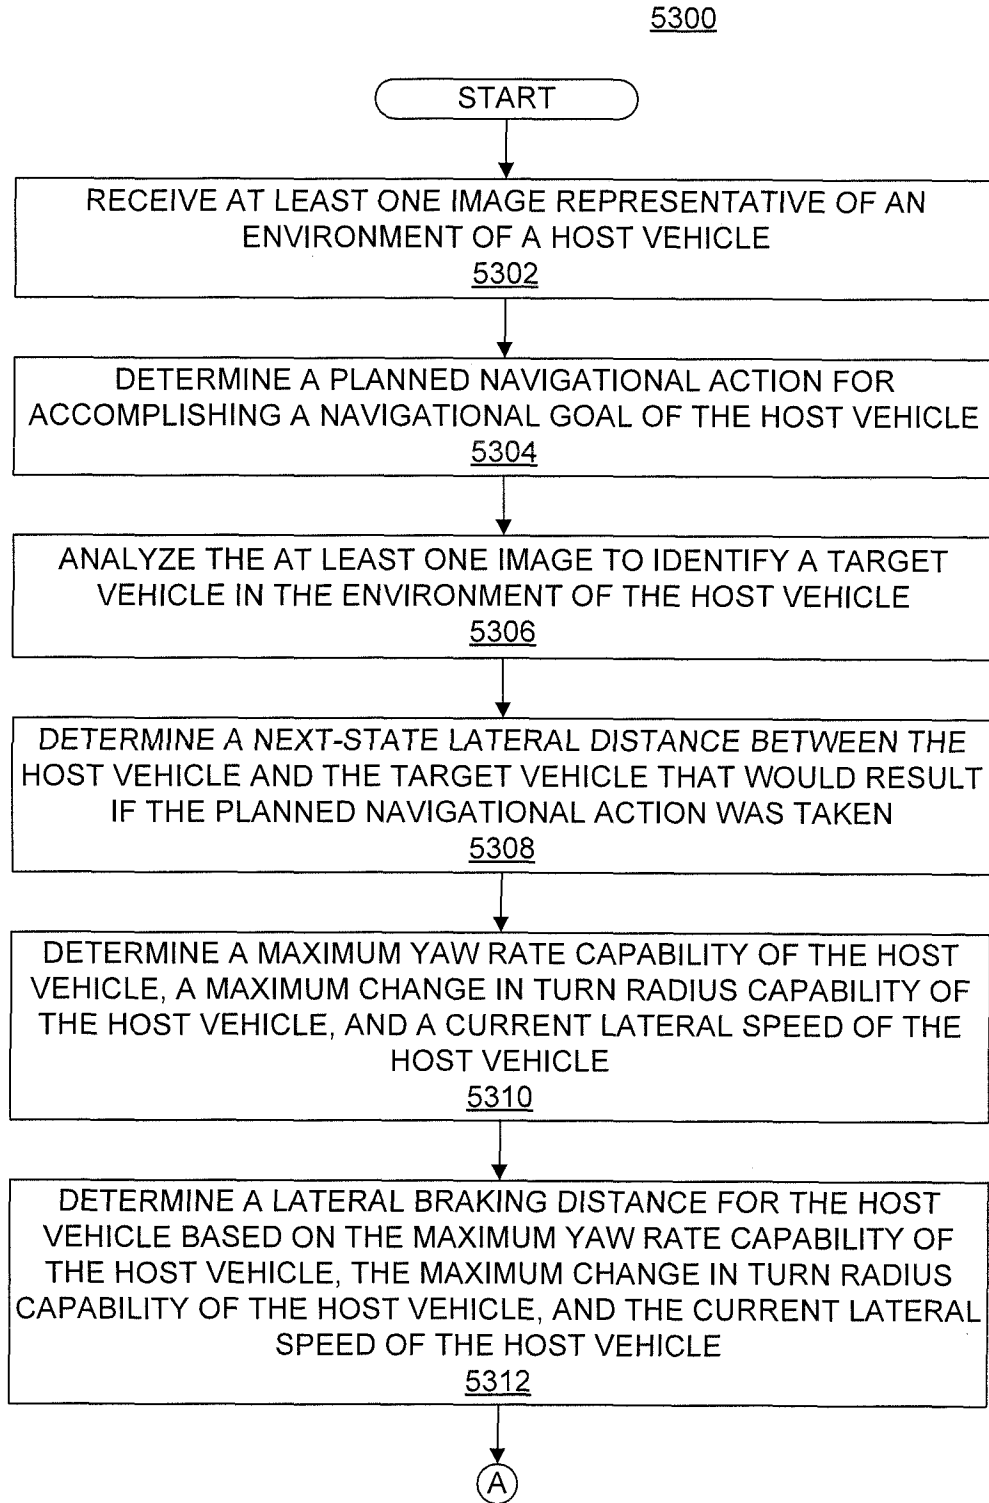


FIG. 53A

5300

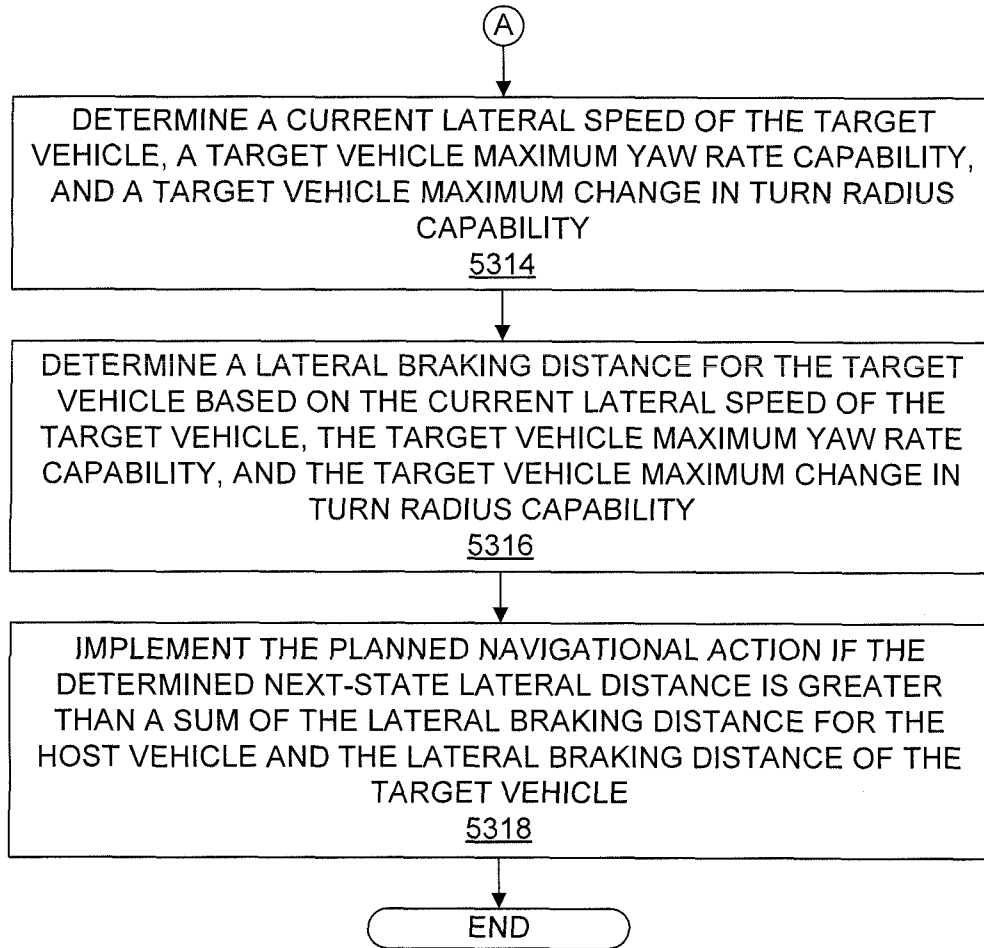


FIG. 53B




FIG. 54




FIG. 55A




FIG. 55B




FIG. 55C


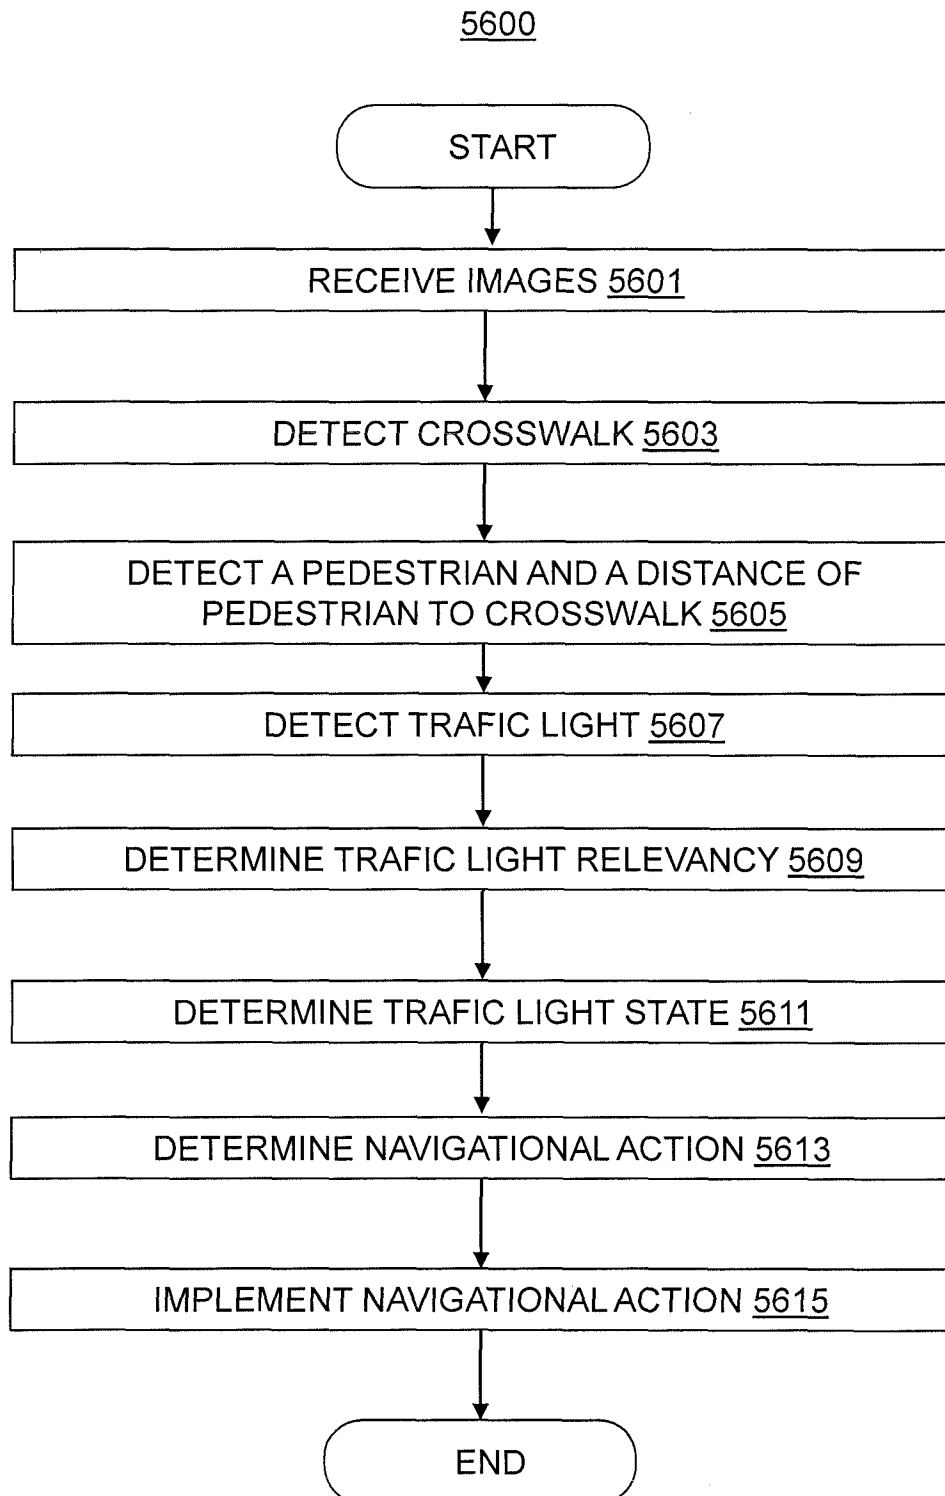


FIG. 55D

**FIG. 56**