



(51) International Patent Classification:

G01C 21/32 (2006.01)

G06K 9/00 (2006.01)

G01C 21/36 (2006.01)

(21) International Application Number:

PCT/US20 19/063 326

(22) International Filing Date:

26 November 2019 (26.11.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/771,335 26 November 2018 (26.11.2018) US

62/795,868 23 January 2019 (23.01.2019) US

62/800,845 04 February 2019 (04.02.2019) US

(71) Applicant: MOBILEYE VISION TECHNOLOGIES

LTD. [IL/IL]; 13 Hartom Street, P.O.B. 45157, Har Hotzvim, 97775 13 Jerusalem (IL).

(72) Inventors; and

(71) Applicants: SHAPIRA, Dori [IL/IL]; Kibbutz Galon, 7955500 Galon (IL). BRAUNSTEIN, Daniel [IL/IL]; 17 Shalmon Street, Apt. 2, 9987500 Tzur Hadassah (IL). GDALYAHU, Yoram [IL/IL]; 4 Kadish Luz St., Apt. 18, 9692813 Jerusalem (IL). VIENTE, Kfir [IL/IL]; 80 Krinitzi Street, 5260109 Ramat Gan (IL). SCHWARTZ, Maxim [IL/IL]; 22 Oren Street, 9079022 Mevaseret Zion (IL). PEKER, Eli [IL/IL]; 29 Ben Maimon St, 9226218 Jerusalem (IL). HUBERMAN, David [IL/IL]; 3 Tzeret Street, 9346303 Jerusalem (IL). LEVIN, Moshe [US/IL]; 7 HaBaron Hirsch Street, 9614807 Jerusalem (IL). BAR-LEV, Jonathan [US/IL]; 41 Zayit Street, 9979700 Karmei Yosef (IL). ZACKAY, Ora [IL/IL]; 26 HaCohananim St., 4501614 Hod HaSharon (IL). HAYAT, Adi [IL/IL]; 10 Haluakh Haivri St, 7173077 Modiin (IL). FERENCZ, Andras [US/US]; 11 Green St, Princeton, New Jersey 08542 (US). EAGELBERG, Dor [IL/IL]; 14 Gundar Avraham St, 4722369 Ramat Hasharon (IL).

(54) Title: LANE MAPPING AND NAVIGATION




FIG. 36

(57) Abstract: Systems and methods are disclosed for mapping lanes for use in vehicle navigation. In one implementation, at least one processing device may be programmed to receive navigational information from a first vehicle and a second vehicle that have navigated along a road segment including a lane split feature; receive at least one image associated with the road segment; determine, from the first navigational information, a first actual trajectory of the first vehicle and a second actual trajectory of the second vehicle; determine a divergence between the first actual trajectory and the second actual trajectory; determine, based on analysis of the at least one image, that the divergence between the first actual trajectory and the second actual trajectory is indicative of the lane split feature; and update a vehicle road navigation model to include a first target trajectory and a second target trajectory that branches from the first target trajectory after the lane split feature.

(74) **Agent: LOMBARDI, Anthony J.;** Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 901 New York Avenue, NW, Washington, District of Columbia 20001-4413 (US).

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

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

Published:

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

LANE MAPPING AND NAVIGATIONCross-References to Related Applications

[0001] This application claims the benefit of priority of United States Provisional Application No. 62/771,335, filed on November 26, 2018; United States Provisional Application No. 62/795,868, filed on January 23, 2019; and United States Provisional Application No. 62/800,845, filed on February 4, 2019. All of the foregoing applications are incorporated herein by reference in their entirety.

BACKGROUNDTechnical Field

[0002] The present disclosure relates generally to autonomous vehicle navigation. In particular, this disclosure relates to systems and methods for mapping lane splits or lane merges and navigating using the mapped lane splits or lane merges.

Background Information

[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) and may also use 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, and travel from one road to another road at appropriate intersections or interchanges. Harnessing and interpreting vast volumes of information collected by an autonomous vehicle as the vehicle travels to its destination poses a multitude of design challenges. The sheer quantity of data (e.g., captured image data, map data, GPS data, sensor data, etc.) that an autonomous vehicle may need to analyze, access, and/or store poses challenges that can in fact limit or even adversely affect autonomous navigation. Furthermore, if an autonomous vehicle relies on traditional mapping technology to navigate, the sheer volume of data needed to store and update the map poses daunting challenges.

[0004] In addition to the collection of data for updating the map, autonomous vehicles must be able to use the map for navigation. Accordingly, the size and detail of the map must be optimized, as well as the construction and transmission thereof.

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.

[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 first navigational information from a first vehicle that has navigated along a road segment, wherein the road segment includes a lane split feature, and wherein the road segment includes at least a first travel lane prior to the lane split feature that transitions into at least a second travel lane and a third travel lane after the lane split feature; receive second navigation information from a second vehicle that has navigated along the road segment; receive at least one image associated with the road segment; determine, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the second travel lane of the road segment; determine, from the second navigational information, a second actual trajectory of the second vehicle along the first travel lane and the third travel lane of the road segment; determine a divergence between the first actual trajectory and the second actual trajectory; determine, based on analysis of the at least one image, that the divergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane split feature; and update a vehicle road navigation model to include: a first target trajectory corresponding to the first travel lane prior to the road split feature and extending along the second travel lane after the lane split feature; and a second target trajectory that branches from the first target trajectory and extends along the third travel lane after the lane split feature.

[0007] In an embodiment, a method for mapping lane splits for use in vehicle navigation is disclosed. The method may comprise receiving first navigational information from a first vehicle that has navigated along a road segment, wherein the road segment includes a lane split feature, and wherein the road segment includes at least a first travel lane prior to the lane split feature that transitions into at least a second travel lane and a third travel lane after the lane split feature; receiving second navigation information from a second vehicle that has navigated along the road segment; receiving at least one image associated with the road segment; determining, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the second travel lane of the road segment; determining, from the second navigational information, a second actual trajectory of the second vehicle along the first travel lane and the third travel lane of the road segment; determining a divergence between the first actual trajectory and the second actual trajectory; determining, based on analysis of the at least one image, that the divergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane split feature; and updating a vehicle road navigation model to include: a first target trajectory corresponding to the first travel lane prior to the road split feature and extending along the second travel lane after the lane split feature; and a second target trajectory that branches from the first target trajectory and extends along the third travel lane after the lane split feature.

[0008] In an embodiment, a system for navigating a host vehicle along a road segment is disclosed. The system may comprise at least one processing device programmed to receive, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory that corresponds with a first lane of travel along the road segment prior to

a lane split feature associated with the road segment and extends along a second lane of travel of the road segment, subsequent to the lane split feature, and wherein the vehicle road navigation model also includes a second target trajectory that branches from the first target trajectory and extends along a third lane of travel of the road segment, subsequent to the lane split feature; receive information indicative of an environment of the host vehicle; determine, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory; and determine a navigational action to cause the host vehicle to navigate along the determined target trajectory.

[0009] In an embodiment, a method for navigating a host vehicle along a road segment is disclosed. The method may comprise receiving, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory that corresponds with a first lane of travel along the road segment prior to a lane split feature associated with the road segment and extends along a second lane of travel of the road segment subsequent to the lane split feature, and wherein the vehicle road navigation model also includes a second target trajectory that branches from the first target trajectory and extends along a third lane of travel of the road segment subsequent to the lane split feature; receiving information indicative of an environment of the host vehicle; determining, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory; and determining a navigational action to cause the host vehicle to navigate along the determined target trajectory.

[0010] In an embodiment, a system for mapping lane merges for use in vehicle navigation may comprise at least one processor programmed to receive first navigational information from a first vehicle that has navigated along a road segment, wherein the road segment includes a lane merge feature, and wherein the road segment includes at least a first travel lane and a second travel lane prior to the lane merge feature that transition into a third travel lane after the merge feature; receive second navigation information from a second vehicle that has navigated along the road segment; receive at least one image associated with the road segment; determine, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the third travel lane of the road segment; determine, from the second navigational information, a second actual trajectory of the second vehicle along the second travel lane and the third travel lane of the road segment; determine a convergence between the first actual trajectory and the second actual trajectory; determine, based on analysis of the at least one image, that the convergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane merge feature; and update a vehicle road navigation model to include: a first target trajectory corresponding to the first travel lane prior to the lane merge feature and extending along the third travel lane after the lane merge feature; and a second target trajectory that extends along the second travel lane before the lane merge feature and joins with the first target trajectory.

[0011] In an embodiment, a system for navigating a host vehicle along a road segment may include at least one processor programmed to receive, from a server-based system, a vehicle road

navigation model, wherein the vehicle road navigation model includes a first target trajectory corresponding to a first travel lane prior to a lane merge feature and extending along a third travel lane after the lane merge feature, and wherein the vehicle road navigation model also includes a second target trajectory that extends along a second travel lane before the lane merge feature and joins with the first target trajectory. The at least one processor may be further programmed to receive information indicative of an environment of the host vehicle, determine, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory, and determine a navigational action to cause the host vehicle to navigate along the determined target trajectory.

[0012] Consistent with other disclosed embodiments, non-transitory computer-readable storage media may store program instructions, which are executed by at least one processing device and perform any of the methods described herein.

[0013] 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

[0014] The accompanying drawings, which are incorporated in and constitute a part of this disclosure, illustrate various disclosed embodiments. In the drawings:

[0015] FIG. 1 is a diagrammatic representation of an exemplary system consistent with the disclosed embodiments.

[0016] FIG. 2A is a diagrammatic side view representation of an exemplary vehicle including a system consistent with the disclosed embodiments.

[0017] FIG. 2B is a diagrammatic top view representation of the vehicle and system shown in FIG. 2A consistent with the disclosed embodiments.

[0018] FIG. 2C is a diagrammatic top view representation of another embodiment of a vehicle including a system consistent with the disclosed embodiments.

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

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

[0021] FIG. 2F is a diagrammatic representation of exemplary vehicle control systems consistent with the disclosed embodiments.

[0022] 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.

[0023] 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.

[0024] FIG. 3C is an illustration of the camera mount shown in FIG. 3B from a different perspective consistent with the disclosed embodiments.

[0025] 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.

5 [0026] 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.

[0027] 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.

[0028] 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.

10 [0029] 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.

[0030] FIG. 5D is a flowchart showing an exemplary process for detecting traffic lights in a set of images consistent with the disclosed embodiments.

15 [0031] 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.

[0032] FIG. 5F is a flowchart showing an exemplary process for determining whether a leading vehicle is changing lanes consistent with the disclosed embodiments.

[0033] 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.

20 [0034] 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 embodiments.

[0035] FIG. 8 shows a sparse map for providing autonomous vehicle navigation, consistent with the disclosed embodiments.

25 [0036] FIG. 9A illustrates a polynomial representation of a portions of a road segment consistent with the disclosed embodiments.

[0037] FIG. 9B illustrates a curve in three-dimensional space representing a target trajectory of a vehicle, for a particular road segment, included in a sparse map consistent with the disclosed embodiments.

30 [0038] FIG. 10 illustrates example landmarks that may be included in sparse map consistent with the disclosed embodiments.

[0039] FIG. 11A shows polynomial representations of trajectories consistent with the disclosed embodiments.

[0040] FIGS. 11B and 11C show target trajectories along a multi-lane road consistent with disclosed embodiments.

35 [0041] FIG. 11D shows an example road signature profile consistent with disclosed embodiments.

[0042] FIG. 12 is a schematic illustration of a system that uses crowd-sourcing data received from a plurality of vehicles for autonomous vehicle navigation, consistent with the disclosed embodiments.

[0043] FIG. 13 illustrates an example autonomous vehicle road navigation model represented by a plurality of three dimensional splines, consistent with the disclosed embodiments.

[0044] FIG. 14 shows a map skeleton generated from combining location information from many drives, consistent with the disclosed embodiments.

[0045] FIG. 15 shows an example of all longitudinal alignment of two drives with example signs as landmarks, consistent with the disclosed embodiments.

[0046] FIG. 16 shows an example of all longitudinal alignment of many drives with an example sign as a landmark, consistent with the disclosed embodiments.

[0047] FIG. 17 is a schematic illustration of a system for generating drive data using a camera, a vehicle, and a server, consistent with the disclosed embodiments.

[0048] FIG. 18 is a schematic illustration of a system for crowdsourcing a sparse map, consistent with the disclosed embodiments.

[0049] FIG. 19 is a flowchart showing an exemplary process for generating a sparse map for autonomous vehicle navigation along a road segment, consistent with the disclosed embodiments.

[0050] FIG. 20 illustrates a block diagram of a server consistent with the disclosed embodiments.

[0051] FIG. 21 illustrates a block diagram of a memory consistent with the disclosed embodiments.

[0052] FIG. 22 illustrates a process of clustering vehicle trajectories associated with vehicles, consistent with the disclosed embodiments.

[0053] FIG. 23 illustrates a navigation system for a vehicle, which may be used for autonomous navigation, consistent with the disclosed embodiments.

[0054] FIG. 24 is a flowchart showing an example process for generating a road navigation model for use in autonomous vehicle navigation, consistent with the disclosed embodiments.

[0055] FIG. 25 illustrates a block diagram of a memory consistent with the disclosed embodiments.

[0056] FIG. 26 is a flowchart showing an example process for autonomously navigating a vehicle along a road segment, consistent with the disclosed embodiments.

[0057] FIG. 27 illustrates a block diagram of a memory consistent with the disclosed embodiments.

[0058] FIG. 28A illustrates an example of drive data from four separate drives, consistent with the disclosed embodiments.

[0059] FIG. 28B illustrates an example of drive data from five separate drives, consistent with the disclosed embodiments.

[0060] FIG. 228BC illustrates an example of vehicle paths determined from drive data from five separate drives, consistent with the disclosed embodiments.

[0061] FIG. 229 is a flowchart showing an example process for determining a line representation of a road surface feature extending along a road segment, consistent with the disclosed embodiments.

5 [0062] FIG. 330 illustrates a block diagram of a memory consistent with the disclosed embodiments.

[0063] FIG. 331 is a flowchart showing an example process for collecting road surface information for a road segment, consistent with the disclosed embodiments.

10 [0064] FIG. 332 illustrates a block diagram of a memory consistent with the disclosed embodiments.

[0065] FIG. 33A illustrates an example of a vehicle traversing a lane without using lane markings.

[0066] FIG. 33B illustrates the example of FIG. 33A after position and heading of the vehicle have drifted.

15 [0067] FIG. 33C illustrates the example of FIG. 33B after position and heading have further drifted and the expected location of a landmark significantly differs from its actual location.

[0068] FIG. 34A illustrates an example of a vehicle traversing a lane without using lane markings, consistent with the disclosed embodiments.

20 [0069] FIG. 34B illustrates the example of FIG. 34A with decreased drift of position and heading, consistent with the disclosed embodiments.

[0070] FIG. 34C illustrates the example of FIG. 34B with the expected location of a landmark significantly aligning with its actual location, consistent with the disclosed embodiments.

[0071] FIG. 35 is a flowchart showing an example process for correcting a position of a vehicle navigating a road segment, consistent with the disclosed embodiments.

25 [0072] FIG. 36 illustrates an example road segment having a lane split, consistent with the disclosed embodiments.

[0073] FIG. 37A illustrates example vehicle trajectories that may be used for identifying a lane split, consistent with the disclosed embodiments.

30 [0074] FIG. 37B illustrates example clustering of vehicle trajectories for determining a lane split, consistent with the disclosed embodiments.

[0075] FIG. 37C illustrates example target trajectories for identifying a lane split, consistent with the disclosed embodiments.

[0076] FIG. 38A illustrates an example process for adjusting the location of a branch point, consistent with the disclosed embodiment.

35 [0077] FIG. 38B illustrates example adjusted target trajectories for identifying a lane split, consistent with the disclosed embodiments.

[0078] FIGS. 39A-39C illustrate example anomalies that may arise in target trajectories, consistent with the disclosed embodiments.

[0079] FIG. 40A illustrates an example road segment including a lane split feature, consistent with the disclosed embodiments.

[0080] FIG. 40B illustrates an example image that may be used by a host vehicle in determining a target trajectory associated with a lane split feature, consistent with the disclosed embodiments.

5 [0081] FIG. 40C illustrates another example road segment including a lane split feature, consistent with the disclosed embodiments.

[0082] FIG. 40D illustrates another example image that may be used by a host vehicle in determining a target trajectory associated with a lane split feature, consistent with the disclosed embodiments.

10 [0083] FIG. 41 is a flowchart depicting an exemplary process 4100 for mapping lane splits for use in vehicle navigation, consistent with the disclosed embodiments.

[0084] FIG. 42 is a flowchart depicting an exemplary process 4200 for mapping lane splits for navigating a host vehicle along a road segment, consistent with the disclosed embodiments.

DETAILED DESCRIPTION

15 [0085] 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 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.

[0086] Autonomous Vehicle Overview

25 [0087] 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 of the vehicle. To be autonomous, a vehicle need not be fully automatic (e.g., fully operation 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), but may leave other aspects to the driver (e.g., braking). In some cases, autonomous vehicles may handle some or all aspects of braking, speed control, and/or steering of the vehicle.

35 [0088] As human drivers typically rely on visual cues and observations to control a vehicle, transportation infrastructures are built accordingly, with lane markings, traffic signs, and traffic lights are all 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, 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 the vehicle is traveling, and the vehicle (as well as other vehicles) may use the information to localize itself on the model.

[0089] In some embodiments in this disclosure, an autonomous vehicle may use information obtained while navigating (e.g., from a camera, GPS device, an accelerometer, a speed sensor, a suspension sensor, etc.). In other embodiments, an autonomous vehicle may use information obtained from past navigations by the vehicle (or by other vehicles) while navigating. In yet other embodiments, an autonomous vehicle may use a combination of information obtained while navigating and information obtained from past navigations. The following sections provide an overview of a system consistent with the disclosed embodiments, following by an overview of a forward-facing imaging system and methods consistent with the system. The sections that follow disclose systems and methods for constructing, using, and updating a sparse map for autonomous vehicle navigation.

[0090] System Overview

[0091] 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), 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 device 110 to image acquisition device 120. For example, data interface 128 may include any wired and/or wireless link or links for transmitting image data acquired by image acquisition device 120 to processing unit 110.

[0092] 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.

[0093] Both applications processor 180 and image processor 190 may include various types of 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), a graphics processing unit (GPU), 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., or GPUs available from manufacturers such as NVIDIA®, ATI®, etc. and may include various architectures (e.g., x86 processor, ARM®, etc.).

[0094] 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.

[0095] 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. For example, processing devices such as field-programmable gate arrays (FPGAs), application-specific integrated circuits (ASICs), and the like may be configured using, for example, one or more hardware description languages (HDLs).

[0096] 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.

[0097] 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 1180 and image processor 1190. 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 1120.

[0098] 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), a graphics processing unit (GPU), 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 GPU may also 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.

[0099] 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 (RAM), read only memory (ROM), 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.

[0100] 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.

[0101] In some embodiments, system 100 may include components such as a speed sensor (e.g., a tachometer, a speedometer) for measuring a speed of vehicle 200 and/or an accelerometer (either single axis or multi-axis) for measuring acceleration of vehicle 200.

[0102] 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, a microphone, a keyboard, a pointer device, a track wheel, a camera, a knob, a button, 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.

[0103] 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, a speaker, a tactile device, and/or any other devices for providing output information to a user.

[0104] 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. Systems and methods of generating such a map are discussed below with reference to FIGS. 8-19.

[0105] 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 to FIGS. 2B-2E, below.

[0106] System 100, or various components thereof, may be incorporated into various different 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, either of image capture devices 122a and 124c of vehicle 200, as shown in FIG. 2A, may be part of an ADAS (Advanced Driver Assistance Systems) imaging set.

[0107] 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 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.

[0108] 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 figures on the front and/or back of vehicle 200, etc..

[0109] 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.

[0110] 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.

[0111] 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.

[0112] In some embodiments, system 100 may upload data according to a "high" privacy level, and under setting as 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 of transmit data, such as, captured images and/or limited location information related to a route.

[0113] Other privacy levels are contemplated. 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 entirely 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.

[0114] 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.

[0115] 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.

[0116] 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.

[0117] 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 the 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.

[0118] 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.

[0119] The first image capture device 122 may acquire a plurality of first images relative to a scene associated with the 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.

[0120] 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.

[0121] 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.

[0122] 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.

[0123] 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 captured 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.

[0124] 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.

[0125] Image capture devices 124 and 126 may acquire a plurality of second and third images relative to a scene associated with the 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.

[0126] 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.

[0127] 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,

122, 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).

[0128] 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.

[0129] 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.

[0130] 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.

[0131] 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.

[0132] 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.

[0133] 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.

[0134] 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.

[0135] 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.

[0136] 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).

[0137] 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.

[0138] 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.

[0139] 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 of image capture devices 124 and/or 126 may be located in a lower half of the field of view of the wider FOV of device 122.

[0140] 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 240.

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.

[0141] 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.

[0142] 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 the shield 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.

[0143] 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.

[0144] 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 sensors 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.

[0145] Forward-Facing Multi-Imaging System

[0146] 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). 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 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 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 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.

[0147] For example, in one embodiment, system 100 may implement a three-camera configuration using image capture devices 122, 124, and 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 to about 60 degrees). In some embodiments, image capture device 126 may act as a main or primary camera. Image capture devices 122, 124, and 126 may be positioned behind rearview mirror 310 and positioned substantially side-by-side (e.g., 66 cm apart). Further, in some embodiments, as

discussed above, one or more of image capture devices 122, 124, and 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 the car on image capture devices 122, 124, and 126.

[0148] 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.

[0149] 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, 124, and 126.

[0150] 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.

[0151] The second processing device may receive images from 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.

[0152] 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.

[0153] 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.

[0154] 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 or more sensors (e.g., radar, lidar, acoustic sensors, information received from one or more transceivers outside of a vehicle, etc.).

[0155] 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.

[0156] 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.

[0157] 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, application processor 180 and/or image processor 190 may execute the instructions stored in any of modules 402, 404, 406, and 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 application 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.

[0158] 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 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.

[0159] 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., 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, 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 of a deep neural network) or an untrained system, such as a system that may be configured to use computer vision algorithms to detect and/or label objects in an environment from which sensory information was captured and processed. In one embodiment, stereo image analysis module 404 and/or other image processing modules may be configured to use a combination of a trained and untrained system.

[0160] 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 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 sensory input (e.g., information from radar) and input from other systems of vehicle 200, such as a 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.

[0156] 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.

[0162] 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.

[0163] 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.

[0164] 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.

[0165] 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. 44. 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 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 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.

[0166] 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 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 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.

[0167] 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 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.

[0168] 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, 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.

[0169] 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.

[0170] 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.

[0171] 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.

[0172] 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.

[0173] 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, 552, 554, and 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.

[0174] 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.

[0175] 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.

[0176] 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.

[0177] 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, 562, and 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.

[0178] 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).

[0179] 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).

[0180] At step 574, processing unit 110 may determine a look-ahead point (expressed in coordinates, $ass(x_i, \Sigma)$) 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 15 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.

5 [0181] 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., $\arctan((x/z))$. 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: $(2 * \text{speed of vehicle 200} / \text{look-ahead distance})$.

15 [0182] 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.

20 [0183] 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.

35 [0184] 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 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: $(\delta x^2 + \delta y^2) / 2 / (\delta x)$, where δx represents the lateral distance traveled and δy represents the longitudinal distance traveled. If the difference between the lateral distance traveled and the expected 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 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.

[0185] 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 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 combination of analyses and weights.

[0186] 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 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 of images via two or more data interfaces. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0187] 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 exits/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.

[0188] 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.

[0189] 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.

[0190] 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. 55A-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. 66, 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.

[0191] 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.”

[0192] 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.

[0193] 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.

[0194] 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.

[0195] Analysis of captured images may allow for the generation and use of a sparse map model for autonomous vehicle navigation. In addition, analysis of captured images may allow for the localization of an autonomous vehicle using identified lane markings. Embodiments for detection of particular characteristics based on one or more particular analyses of captured images and for navigation of an autonomous vehicle using a sparse map model will be discussed below with reference to FIGs. 8-28.

[0196] Sparse Road Model for Autonomous Vehicle Navigation

[0197] In some embodiments, the disclosed systems and methods may use a sparse map for autonomous vehicle navigation. In particular, the sparse map may be for autonomous vehicle navigation along a road segment. For example, the sparse map may provide sufficient information for navigating an autonomous vehicle without storing and/or updating a large quantity of data. As discussed below in further detail, an autonomous vehicle may use the sparse map to navigate one or more roads based on one or more stored trajectories.

[0198] Sparse Map for Autonomous Vehicle Navigation

[0199] In some embodiments, the disclosed systems and methods may generate a sparse map for autonomous vehicle navigation. For example, the sparse map may provide sufficient information for navigation without requiring excessive data storage or data transfer rates. As discussed below in further detail, a vehicle, (which may be an autonomous vehicle) may use the sparse map to navigate one or more roads. For example, in some embodiments, the sparse map may include data related to a road and potentially landmarks along the road that may be sufficient for vehicle navigation, but which also exhibit small data footprints. For example, the sparse data maps described in detail below may require significantly less storage space and data transfer bandwidth as compared with digital maps including detailed map information, such as image data collected along a road.

[0200] For example, rather than storing detailed representations of a road segment, the sparse data map may store three-dimensional polynomial representations of preferred vehicle paths along a road. These paths may require very little data storage space. Further, in the described sparse data maps, landmarks may be identified and included in the sparse map road model to aid in navigation. These landmarks may be located at any spacing suitable for enabling vehicle navigation, but in some cases, such landmarks need not be identified and included in the model at high densities and short spacings. Rather, in some cases, navigation may be possible based on landmarks that are spaced apart by at least 500 meters,

at least 100 meters, at least 500 meters, at least 1 kilometer, or at least 2 kilometers. As will be discussed in more detail in other sections, the sparse map may be generated based on data collected or measured by vehicles equipped with various sensors and devices, such as image capture devices, Global Positioning System sensors, motion sensors, etc., as the vehicles travel along roadways. In some cases, the sparse map may be generated based on data collected during multiple drives of one or more vehicles along a particular roadway. Generating a sparse map using multiple drives of one or more vehicles may be referred to as "crowdsourcing" a sparse map.

[0201] Consistent with disclosed embodiments, an autonomous vehicle system may use a sparse map for navigation. For example, the disclosed systems and methods may distribute a sparse map for generating a road navigation model for an autonomous vehicle and may navigate an autonomous vehicle along a road segment using a sparse map and/or a generated road navigation model. Sparse maps consistent with the present disclosure may include one or more three-dimensional contours that may represent predetermined trajectories that autonomous vehicles may traverse as they move along associated road segments.

[0202] Sparse maps consistent with the present disclosure may also include data representing one or more road features. Such road features may include recognized landmarks, road signature profiles, and any other road-related features useful in navigating a vehicle. Sparse maps consistent with the present disclosure may enable autonomous navigation of a vehicle based on relatively small amounts of data included in the sparse map. For example, rather than including detailed representations of a road, such as road edges, road curvature, images associated with road segments, or data detailing other physical features associated with a road segment, the disclosed embodiments of the sparse map may require relatively little storage space (and relatively little bandwidth when portions of the sparse map are transferred to a vehicle), but may still adequately provide for autonomous vehicle navigation. The small data footprint of the disclosed sparse maps, discussed in further detail below, may be achieved in some embodiments by storing representations of road-related elements that require small amounts of data but still enable autonomous navigation.

[0203] For example, rather than storing detailed representations of various aspects of a road, the disclosed sparse maps may store polynomial representations of one or more trajectories that a vehicle may follow along the road. Thus, rather than storing (or having to transfer) details regarding the physical nature of the road to enable navigation along the road, using the disclosed sparse maps, a vehicle may be navigated along a particular road segment without, in some cases, having to interpret physical aspects of the road, but rather, by aligning its path of travel with a trajectory (e.g., a polynomial spline) along the particular road segment. In this way, the vehicle may be navigated based mainly upon the stored trajectory (e.g., a polynomial spline) that may require much less storage space than an approach involving storage of roadway images, road parameters, road layout, etc.

[0204] In addition to the stored polynomial representations of trajectories along a road segment, the disclosed sparse maps may also include small data objects that may represent a road feature. In some embodiments, the small data objects may include digital signatures, which are derived from a digital

image (or a digital signal) that was obtained by a sensor (e.g., a camera or other sensor, such as a suspension sensor), onboard a vehicle, traveling along the road segment. The digital signature may have a reduced size relative to the signal that was acquired by the sensor. In some embodiments, the digital signature may be created to be compatible with a classifier function that is configured to detect and to identify the road feature from the signal that is acquired by the sensor, for example, during a subsequent drive. In some embodiments, a digital signature may be created such that the digital signature has a footprint that is as small as possible, while retaining the ability to correlate or match the road feature with the stored signature based on an image (or a digital signal generated by a sensor, if the stored signature is not based on an image and/or includes other data) of the road feature that is captured by a camera onboard a vehicle traveling along the same road segment at a subsequent time.

[0205] In some embodiments, a size of the data objects may be further associated with a uniqueness of the road feature. For example, for a road feature that is detectable by a camera onboard a vehicle, and where the camera system onboard the vehicle is coupled to a classifier that is capable of distinguishing the image data corresponding to that road feature as being associated with a particular type of road feature, for example, a road sign, and where such a road sign is locally unique in that area (e.g., there is no identical road sign or road sign of the same type nearby), it may be sufficient to store data indicating the type of the road feature and its location.

[0206] As will be discussed in further detail below, road features (e.g., landmarks along a road segment) may be stored as small data objects that may represent a road feature in relatively few bytes, while at the same time providing sufficient information for recognizing and using such a feature for navigation. In one example, a road sign may be identified as a recognized landmark on which navigation of a vehicle may be based. A representation of the road sign may be stored in the sparse map to include, e.g., a few bytes of data indicating a type of landmark (e.g., a stop sign) and a few bytes of data indicating a location of the landmark (e.g., coordinates). Navigating based on such data-light representations of the landmarks (e.g., using representations sufficient for locating, recognizing, and navigating based upon the landmarks) may provide a desired level of navigational functionality associated with sparse maps without significantly increasing the data overhead associated with the sparse maps. This lean representation of landmarks (and other road features) may take advantage of the sensors and processors included onboard such vehicles that are configured to detect, identify, and/or classify certain road features.

[0207] When, for example, a sign or even a particular type of a sign is locally unique (e.g., when there is no other sign or no other sign of the same type) in a given area, the sparse map may use data indicating a type of a landmark (a sign or a specific type of sign), and during navigation (e.g., autonomous navigation) when a camera onboard an autonomous vehicle captures an image of the area including a sign (or of a specific type of sign), the processor may process the image, detect the sign (if indeed present in the image), classify the image as a sign (or as a specific type of sign), and correlate the location of the image with the location of the sign as stored in the sparse map.

[0208] Generating a Sparse Map

[0209] In some embodiments, a sparse map may include at least one line representation of a road surface feature extending along a road segment and a plurality of landmarks associated with the road segment. In certain aspects, the sparse map may be generated via "crowdsourcing," for example, through image analysis of a plurality of images acquired as one or more vehicles traverse the road segment.

[0210] FIG. 8 shows a sparse map 800 that one or more vehicles, e.g., vehicle 200 (which may be an autonomous vehicle), may access for providing autonomous vehicle navigation. Sparse map 800 may be stored in a memory, such as memory 114 or 150. Such memory devices may include any types of non-transitory storage devices or computer-readable media. For example, in some embodiments, memory 114 or 150 may include hard drives, compact discs, flash memory, magnetic based memory devices, optical based memory devices, etc. In some embodiments, sparse map 800 may be stored in a database (e.g., map database 160) that may be stored in memory 114 or 150, or other types of storage devices.

[0211] In some embodiments, sparse map 800 may be stored on a storage device or a non-transitory computer-readable medium provided onboard vehicle 200 (e.g., a storage device included in a navigation system onboard vehicle 200). A processor (e.g., processing unit 110) provided on vehicle 200 may access sparse map 800 stored in the storage device or computer-readable medium provided onboard vehicle 200 in order to generate navigational instructions for guiding the autonomous vehicle 200 as the vehicle traverses a road segment.

[0212] Sparse map 800 need not be stored locally with respect to a vehicle, however. In some embodiments, sparse map 800 may be stored on a storage device or computer-readable medium provided on a remote server that communicates with vehicle 200 or a device associated with vehicle 200. A processor (e.g., processing unit 110) provided on vehicle 200 may receive data included in sparse map 800 from the remote server and may execute the data for guiding the autonomous driving of vehicle 200. In such embodiments, the remote server may store all of sparse map 800 or only a portion thereof. Accordingly, the storage device or computer-readable medium provided onboard vehicle 200 and/or onboard one or more additional vehicles may store the remaining portion(s) of sparse map 800.

[0213] Furthermore, in such embodiments, sparse map 800 may be made accessible to a plurality of vehicles traversing various road segments (e.g., tens, hundreds, thousands, or millions of vehicles, etc.). It should be noted also that sparse map 800 may include multiple sub-maps. For example, in some embodiments, sparse map 800 may include hundreds, thousands, millions, or more, of sub-maps that may be used in navigating a vehicle. Such sub-maps may be referred to as local maps, and a vehicle traveling along a roadway may access any number of local maps relevant to a location in which the vehicle is traveling. The local map sections of sparse map 800 may be stored with a Global Navigation Satellite System (GNSS) key as an index to the database of sparse map 800. Thus, while computation of steering angles for navigating a host vehicle in the present system may be performed without reliance upon a GNSS position of the host vehicle, road features, or landmarks, such GNSS information may be used for retrieval of relevant local maps.

[0214] Collection of data and generation of sparse map 800 is covered in greater detail below, for example, with respect to FIG. 19. In general, however, sparse map 800 may be generated based on

data collected from one or more vehicles as they travel along roadways. For example, using sensors aboard the one or more vehicles (e.g., cameras, speedometers, GPS, accelerometers, etc.), the trajectories that the one or more vehicles travel along a roadway may be recorded, and the polynomial representation of a preferred trajectory for vehicles making subsequent trips along the roadway may be determined based on the collected trajectories travelled by the one or more vehicles. Similarly, data collected by the one or more vehicles may aid in identifying potential landmarks along a particular roadway. Data collected from traversing vehicles may also be used to identify road profile information, such as road width profiles, road roughness profiles, traffic line spacing profiles, road conditions, etc. Using the collected information, sparse map 800 may be generated and distributed (e.g., for local storage or via on-the-fly data transmission) for use in navigating one or more autonomous vehicles. However, in some embodiments, map generation may not end upon initial generation of the map. As will be discussed in greater detail below, sparse map 800 may be continuously or periodically updated based on data collected from vehicles as those vehicles continue to traverse roadways included in sparse map 800.

[0215] Data recorded in sparse map 800 may include position information based on Global Positioning System (GPS) data. For example, location information may be included in sparse map 800 for various map elements, including, for example, landmark locations, road profile locations, etc. Locations for map elements included in sparse map 800 may be obtained using GPS data collected from vehicles traversing a roadway. For example, a vehicle passing an identified landmark may determine a location of the identified landmark using GPS position information associated with the vehicle and a determination of a location of the identified landmark relative to the vehicle (e.g., based on image analysis of data collected from one or more cameras on board the vehicle). Such location determinations of an identified landmark (or any other feature included in sparse map 800) may be repeated as additional vehicles pass the location of the identified landmark. Some or all of the additional location determinations may be used to refine the location information stored in sparse map 800 relative to the identified landmark. For example, in some embodiments, multiple position measurements relative to a particular feature stored in sparse map 800 may be averaged together. Any other mathematical operations, however, may also be used to refine a stored location of a map element based on a plurality of determined locations for the map element.

[0216] The sparse map of the disclosed embodiments may enable autonomous navigation of a vehicle using relatively small amounts of stored data. In some embodiments, sparse map 800 may have a data density (e.g., including data representing the target trajectories, landmarks, and any other stored road features) of less than 2 MB per kilometer of roads, less than 1 MB per kilometer of roads, less than 500 kB per kilometer of roads, or less than 100 kB per kilometer of roads. In some embodiments, the data density of sparse map 800 may be less than 10 kB per kilometer of roads or even less than 2 kB per kilometer of roads (e.g., 1.65 kB per kilometer), or no more than 10 kB per kilometer of roads, or no more than 20 kB per kilometer of roads. In some embodiments, most, if not all, of the roadways of the United States may be navigated autonomously using a sparse map having a total of 4 GB or less of data. These

data density values may represent an average over an entire sparse map 800, over a local map within sparse map 800, and/or over a particular road segment within sparse map 800.

[0217] As noted, sparse map 800 may include representations of a plurality of target trajectories 810 for guiding autonomous driving or navigation along a road segment. Such target trajectories may be stored as three-dimensional splines. The target trajectories stored in sparse map 800 may be determined based on two or more reconstructed trajectories of prior traversals of vehicles along a particular road segment, for example, as discussed with respect to FIG. 229. A road segment may be associated with a single target trajectory or multiple target trajectories. For example, on a two-lane road, a first target trajectory may be stored to represent an intended path of travel along the road in a first direction, and a second target trajectory may be stored to represent an intended path of travel along the road in another direction (e.g., opposite to the first direction). Additional target trajectories may be stored with respect to a particular road segment. For example, on a multi-lane road one or more target trajectories may be stored representing intended paths of travel for vehicles in one or more lanes associated with the multi-lane road. In some embodiments, each lane of a multi-lane road may be associated with its own target trajectory. In other embodiments, there may be fewer target trajectories stored than lanes present on a multi-lane road. In such cases, a vehicle navigating the multi-lane road may use any of the stored target trajectories to guide its navigation by taking into account an amount of lane offset from a lane for which a target trajectory is stored (e.g., if a vehicle is traveling in the left most lane of a three lane highway, and a target trajectory is stored only for the middle lane of the highway, the vehicle may navigate using the target trajectory of the middle lane by accounting for the amount of lane offset between the middle lane and the left-most lane when generating navigational instructions).

[0218] In some embodiments, the target trajectory may represent an ideal path that a vehicle should take as the vehicle travels. The target trajectory may be located, for example, at an approximate center of a lane of travel. In other cases, the target trajectory may be located elsewhere relative to a road segment. For example, a target trajectory may approximately coincide with a center of a road, an edge of a road, or an edge of a lane, etc. In such cases, navigation based on the target trajectory may include a determined amount of offset to be maintained relative to the location of the target trajectory. Moreover, in some embodiments, the determined amount of offset to be maintained relative to the location of the target trajectory may differ based on a type of vehicle (e.g., a passenger vehicle including two axles may have a different offset from a truck including more than two axles along at least a portion of the target trajectory).

[0219] Sparse map 800 may also include data relating to a plurality of predetermined landmarks 820 associated with particular road segments, local maps, etc. As discussed in greater detail below, these landmarks may be used in navigation of the autonomous vehicle. For example, in some embodiments, the landmarks may be used to determine a current position of the vehicle relative to a stored target trajectory. With this position information, the autonomous vehicle may be able to adjust a heading direction to match a direction of the target trajectory at the determined location.

[0220] The plurality of landmarks may be identified and stored in sparse map 800 at any suitable spacing. In some embodiments, landmarks may be stored at relatively high densities (e.g., every few meters or more). In some embodiments, however, significantly larger landmark spacing values may be employed. For example, in sparse map 800, identified (or recognized) landmarks may be spaced apart by 10 meters, 20 meters, 50 meters, 100 meters, 1 kilometer, or 2 kilometers. In some cases, the identified landmarks may be located at distances of even more than 2 kilometers apart.

[0221] Between landmarks, and therefore between determinations of vehicle position relative to a target trajectory, the vehicle may navigate based on dead reckoning in which the vehicle uses sensors to determine its ego motion and estimate its position relative to the target trajectory. Because errors may accumulate during navigation by dead reckoning, over time the position determinations relative to the target trajectory may become increasingly less accurate. The vehicle may use landmarks occurring in sparse map 800 (and their known locations) to remove the dead reckoning-induced errors in position determination. In this way, the identified landmarks included in sparse map 800 may serve as navigational anchors from which an accurate position of the vehicle relative to a target trajectory may be determined. Because a certain amount of error may be acceptable in position location, an identified landmark need not always be available to an autonomous vehicle. Rather, suitable navigation may be possible even based on landmark spacings, as noted above, of 10 meters, 20 meters, 50 meters, 100 meters, 500 meters, 1 kilometer, 2 kilometers, or more. In some embodiments, a density of 1 identified landmark every 1 km of road may be sufficient to maintain a longitudinal position determination accuracy within 1 m. Thus, not every potential landmark appearing along a road segment need be stored in sparse map 800.

[0222] Moreover, in some embodiments, lane markings may be used for localization of the vehicle during landmark spacings. By using lane markings during landmark spacings, the accumulation of during navigation by dead reckoning may be minimized. In particular, such localization is discussed below with respect to FIG. 35.

[0223] In addition to target trajectories and identified landmarks, sparse map 800 may include information relating to various other road features. For example, FIG. 9A illustrates a representation of curves along a particular road segment that may be stored in sparse map 800. In some embodiments, a single lane of a road may be modeled by a three-dimensional polynomial description of left and right sides of the road. Such polynomials representing left and right sides of a single lane are shown in FIG. 9A. Regardless of how many lanes a road may have, the road may be represented using polynomials in a way similar to that illustrated in FIG. 9A. For example, left and right sides of a multi-lane road may be represented by polynomials similar to those shown in FIG. 9A, and intermediate lane markings included on a multi-lane road (e.g., dashed markings representing lane boundaries, solid yellow lines representing boundaries between lanes traveling in different directions, etc.) may also be represented using polynomials such as those shown in FIG. 9A.

[0224] As shown in FIG. 9A, a lane 900 may be represented using polynomials (e.g., a first order, second order, third order, or any suitable order polynomials). For illustration, lane 900 is shown as

a two-dimensional lane and the polynomials are shown as two-dimensional polynomials. As depicted in FIG. 9A, lane 900 includes a left side 910 and a right side 920. In some embodiments, more than one polynomial may be used to represent a location of each side of the road or lane boundary. For example, each of left side 910 and right side 920 may be represented by a plurality of polynomials of any suitable length. In some cases, the polynomials may have a length of about 100 m, although other lengths greater than or less than 100 m may also be used. Additionally, the polynomials can overlap with one another in order to facilitate seamless transitions in navigating based on subsequently encountered polynomials as a host vehicle travels along a roadway. For example, each of left side 910 and right side 920 may be represented by a plurality of third order polynomials separated into segments of about 100 meters in length (an example of the first predetermined range), and overlapping each other by about 50 meters. The polynomials representing the left side 910 and the right side 920 may or may not have the same order. For example, in some embodiments, some polynomials may be second order polynomials, some may be third order polynomials, and some may be fourth order polynomials.

[0225] In the example shown in FIG. 9A, left side 910 of lane 900 is represented by two groups of third order polynomials. The first group includes polynomial segments 911, 912, and 913. The second group includes polynomial segments 914, 915, and 916. The two groups, while substantially parallel to each other, follow the locations of their respective sides of the road. Polynomial segments 911, 912, 913, 914, 915, and 916 have a length of about 100 meters and overlap adjacent segments in the series by about 50 meters. As noted previously, however, polynomials of different lengths and different overlap amounts may also be used. For example, the polynomials may have lengths of 500 m, 1 km, or more, and the overlap amount may vary from 0 to 50 m, 50 m to 100 m, or greater than 100 m. Additionally, while FIG. 9A is shown as representing polynomials extending in 2D space (e.g., on the surface of the paper), it is to be understood that these polynomials may represent curves extending in three dimensions (e.g., including a height component) to represent elevation changes in a road segment in addition to X-Y curvature. In the example shown in FIG. 9A, right side 920 of lane 900 is further represented by a first group having polynomial segments 921, 922, and 923 and a second group having polynomial segments 924, 925, and 926.

[0226] Returning to the target trajectories of sparse map 800, FIG. 9B shows a three-dimensional polynomial representing a target trajectory for a vehicle traveling along a particular road segment. The target trajectory represents not only the X-Y path that a host vehicle should travel along a particular road segment, but also the elevation change that the host vehicle will experience when traveling along the road segment. Thus, each target trajectory in sparse map 800 may be represented by one or more three-dimensional polynomials, like the three-dimensional polynomial 1950 shown in FIG. 9B. Sparse map 800 may include a plurality of trajectories (e.g., millions or billions or more to represent trajectories of vehicles along various road segments along roadways throughout the world). In some embodiments, each target trajectory may correspond to a spline connecting three-dimensional polynomial segments.

[0227] Regarding the data footprint of polynomial curves stored in sparse map 800, in some embodiments, each third degree polynomial may be represented by four parameters, each requiring four

bytes of data. Suitable representations may be obtained with third degree polynomials requiring about 192 bytes of data for every 100 m. This may translate to approximately 200 KB per hour in data usage/transfer requirements for a host vehicle traveling approximately 100 km/hr.

[0228] Sparse map 800 may describe the lanes network using a combination of geometry descriptors and meta-data. The geometry may be described by polynomials or splines as described above. The meta-data may describe the number of lanes, special characteristics (such as a car-pool lane), and possibly other sparse labels. The total footprint of such indicators may be negligible.

[0229] Accordingly, a sparse map according to embodiments of the present disclosure may include at least one line representation of a road surface feature extending along the road segment, each line representation representing a path along the road segment substantially corresponding with the road surface feature. In some embodiments, as discussed above, the at least one line representation of the road surface feature may include a spline, a polynomial representation, or a curve. Furthermore, in some embodiments, the road surface feature may include at least one of a road edge or a lane marking. Moreover, as discussed below with respect to "crowdsourcing," the road surface feature may be identified through image analysis of a plurality of images acquired as one or more vehicles traverse the road segment.

[0230] As previously noted, sparse map 800 may include a plurality of predetermined landmarks associated with a road segment. Rather than storing actual images of the landmarks and relying, for example, on image recognition analysis based on captured images and stored images, each landmark in sparse map 800 may be represented and recognized using less data than a stored, actual image would require. Data representing landmarks may still include sufficient information for describing or identifying the landmarks along a road. Storing data describing characteristics of landmarks, rather than the actual images of landmarks, may reduce the size of sparse map 800.

[0231] FIG. 10 illustrates examples of types of landmarks that may be represented in sparse map 800. The landmarks may include any visible and identifiable objects along a road segment. The landmarks may be selected such that they are fixed and do not change often with respect to their locations and/or content. The landmarks included in sparse map 800 may be useful in determining a location of vehicle 200 with respect to a target trajectory as the vehicle traverses a particular road segment. Examples of landmarks may include traffic signs, directional signs, general signs (e.g., rectangular signs), roadside fixtures (e.g., lampposts, reflectors, etc.), and any other suitable category. In some embodiments, lane marks on the road may also be included as landmarks in sparse map 800.

[0232] Examples of landmarks shown in FIG. 10 include traffic signs, directional signs, roadside fixtures, and general signs. Traffic signs may include, for example, speed limit signs (e.g., speed limit sign 1000), yield signs (e.g., yield sign 1005), route number signs (e.g., route number sign 1010), traffic light signs (e.g., traffic light sign 1015), stop signs (e.g., stop sign 1020). Directional signs may include a sign that includes one or more arrows indicating one or more directions to different places. For example, directional signs may include a highway sign 1025 having arrows for directing vehicles to different roads

or places, an exit sign 1030 having an arrow directing vehicles off a road, etc. Accordingly, at least one of the plurality of landmarks may include a road sign.

[0233] General signs may be unrelated to traffic. For example, general signs may include billboards used for advertisement, or a welcome board adjacent a border between two countries, states, counties, cities, or towns. FIG. 100 shows a general sign 1040 ("Joe's Restaurant"). Although general sign 1040 may have a rectangular shape, as shown in FIG. 100, general sign 1040 may have other shapes, such as square, circle, triangle, etc.

[0234] Landmarks may also include roadside fixtures. Roadside fixtures may be objects that are not signs, and may not be related to traffic or directions. For example, roadside fixtures may include lampposts (e.g., lamppost 11035), power line posts, traffic light posts, etc.

[0235] Landmarks may also include beacons that may be specifically designed for usage in an autonomous vehicle navigation system. For example, such beacons may include stand-alone structures placed at predetermined intervals to aid in navigating a host vehicle. Such beacons may also include visual/graphical information added to existing road signs (e.g., icons, emblems, bar codes, etc.) that may be identified or recognized by a vehicle traveling along a road segment. Such beacons may also include electronic components. In such embodiments, electronic beacons (e.g., RFID tags, etc.) may be used to transmit non-visual information to a host vehicle. Such information may include, for example, landmark identification and/or landmark location information that a host vehicle may use in determining its position along a target trajectory.

[0236] In some embodiments, the landmarks included in sparse map 800 may be represented by a data object of a predetermined size. The data representing a landmark may include any suitable parameters for identifying a particular landmark. For example, in some embodiments, landmarks stored in sparse map 800 may include parameters such as a physical size of the landmark (e.g., to support estimation of distance to the landmark based on a known size/scale), a distance to a previous landmark, lateral offset, height, a type code (e.g., a landmark type—what type of directional sign, traffic sign, etc.), a GPS coordinate (e.g., to support global localization), and any other suitable parameters. Each parameter may be associated with a data size. For example, a landmark size may be stored using 8 bytes of data. A distance to a previous landmark, a lateral offset, and height may be specified using 12 bytes of data. A type code associated with a landmark such as a directional sign or a traffic sign may require about 2 bytes of data. For general signs, an image signature enabling identification of the general sign may be stored using 50 bytes of data storage. The landmark GPS position may be associated with 6 bytes of data storage. These data sizes for each parameter are examples only, and other data sizes may also be used.

[0237] Representing landmarks in sparse map 800 in this manner may offer a lean solution for efficiently representing landmarks in the database. In some embodiments, signs may be referred to as semantic signs and non-semantic signs. A semantic sign may include any class of signs for which there is a standardized meaning (e.g., speed limit signs, warning signs, directional signs, etc.). A non-semantic sign may include any sign that is not associated with a standardized meaning (e.g., general advertising signs, signs identifying business establishments, etc.). For example, each semantic sign may be

represented with 38 bytes of data (e.g., 8 bytes for size, 12 bytes for distance to previous landmark, lateral offset, and height; 2 bytes for a type code; and 16 bytes for GPS coordinates). Sparse map 800 may use a tag system to represent landmark types. In some cases, each traffic sign or directional sign may be associated with its own tag, which may be stored in the database as part of the landmark identification. For example, the database may include on the order of 1000 different tags to represent various traffic signs, and on the order of about 10000 different tags to represent directional signs. Of course, any suitable number of tags may be used, and additional tags may be created as needed. General purpose signs may be represented in some embodiments using less than about 100 bytes (e.g., about 86 bytes including 8 bytes for size, 12 bytes for distance to previous landmark, lateral offset, and height, 50 bytes for an image signature, and 16 bytes for GPS coordinates).

[0238] Thus, for semantic road signs not requiring an image signature, the data density impact to sparse map 800, even at relatively high landmark densities of about 11 per 50 m, may be on the order of about 760 bytes per kilometer (e.g., 20 landmarks per km \times 38 bytes per landmark = 760 bytes). Even for general purpose signs including an image signature component, the data density impact is about 1.72 kB per km (e.g., 20 landmarks per km \times 86 bytes per landmark = 1,720 bytes). For semantic road signs, this equates to about 76 kB per hour of data usage for a vehicle traveling 100 km/hr. For general purpose signs, this equates to about 170 kB per hour for a vehicle traveling 100 km/hr.

[0239] In some embodiments, a generally rectangular object, such as a rectangular sign, may be represented in sparse map 800 by no more than 100 byte of data. The representation of the generally rectangular object (e.g., general sign 1040) in sparse map 800 may include a condensed image signature (e.g., condensed image signature 1045) associated with the generally rectangular object. This condensed image signature may be used, for example, to aid in identification of a general purpose sign, for example, as a recognized landmark. Such a condensed image signature (e.g., image information derived from actual image data representing an object) may avoid a need for storage of an actual image of an object or a need for comparative image analysis performed on actual images in order to recognize landmarks.

[0240] Referring to FIG. 10, sparse map 800 may include or store a condensed image signature 1045 associated with a general sign 1040, rather than an actual image of general sign 1040. For example, after an image capture device (e.g., image capture device 122, 124, or 126) captures an image of general sign 1040, a processor (e.g., image processor 190) or any other processor that can process images either aboard or remotely located relative to a host vehicle may perform an image analysis to extract/create condensed image signature 1045 that includes a unique signature or pattern associated with general sign 1040. In one embodiment, condensed image signature 1045 may include a shape, color pattern, a brightness pattern, or any other feature that may be extracted from the image of general sign 1040 for describing general sign 1040.

[0241] For example, in FIG. 10, the circles, triangles, and stars shown in condensed image signature 1045 may represent areas of different colors. The pattern represented by the circles, triangles, and stars may be stored in sparse map 800, e.g., within the 50 bytes designated to include an image signature. Notably, the circles, triangles, and stars are not necessarily meant to indicate that such shapes

are stored as part of the image signature. Rather, these shapes are meant to conceptually represent recognizable areas having discernible color differences, textual areas, graphical shapes, or other variations in characteristics that may be associated with a general purpose sign. Such condensed image signatures can be used to identify a landmark in the form of a general sign. For example, the condensed image signature can be used to perform a same-not-same analysis based on a comparison of a stored condensed image signature with image data captured, for example, using a camera onboard an autonomous vehicle.

[0242] Accordingly, the plurality of landmarks may be identified through image analysis of the plurality of images acquired as one or more vehicles traverse the road segment. As explained below with respect to “crowdsourcing,” in some embodiments, the image analysis to identify the plurality of landmarks may include accepting potential landmarks when a ratio of images in which the landmark does appear to images in which the landmark does not appear exceeds a threshold. Furthermore, in some embodiments, the image analysis to identify the plurality of landmarks may include rejecting potential landmarks when a ratio of images in which the landmark does not appear to images in which the landmark does appear exceeds a threshold.

[0243] Returning to the target trajectories a host vehicle may use to navigate a particular road segment, FIG. 11A shows polynomial representations trajectories capturing during a process of building or maintaining sparse map 800. A polynomial representation of a target trajectory included in sparse map 800 may be determined based on two or more reconstructed trajectories of prior traversals of vehicles along the same road segment. In some embodiments, the polynomial representation of the target trajectory included in sparse map 800 may be an aggregation of two or more reconstructed trajectories of prior traversals of vehicles along the same road segment. In some embodiments, the polynomial representation of the target trajectory included in sparse map 800 may be an average of the two or more reconstructed trajectories of prior traversals of vehicles along the same road segment. Other mathematical operations may also be used to construct a target trajectory along a road path based on reconstructed trajectories collected from vehicles traversing along a road segment.

[0244] As shown in FIG. 11A, a road segment 1100 may be travelled by a number of vehicles 200 at different times. Each vehicle 200 may collect data relating to a path that the vehicle took along the road segment. The path traveled by a particular vehicle may be determined based on camera data, accelerometer information, speed sensor information, and/or GPS information, among other potential sources. Such data may be used to reconstruct trajectories of vehicles traveling along the road segment, and based on these reconstructed trajectories, a target trajectory (or multiple target trajectories) may be determined for the particular road segment. Such target trajectories may represent a preferred path of a host vehicle (e.g., guided by an autonomous navigation system) as the vehicle travels along the road segment.

[0245] In the example shown in FIG. 11A, a first reconstructed trajectory 1101 may be determined based on data received from a first vehicle traversing road segment 1100 at a first time period (e.g., day 1), a second reconstructed trajectory 1102 may be obtained from a second vehicle traversing road segment 1100 at a second time period (e.g., day 2), and a third reconstructed trajectory 1103 may be

obtained from a third vehicle traversing road segment 1100 at a third time period (e.g., 1 day³). Each trajectory 1101, 1102, and 1103 may be represented by a polynomial, such as a three-dimensional polynomial. It should be noted that in some embodiments, any of the reconstructed trajectories may be assembled onboard the vehicles traversing road segment 1100.

[0246] Additionally, or alternatively, such reconstructed trajectories may be determined on a server side based on information received from vehicles traversing road segment 1100. For example, in some embodiments, vehicles 2200 may transmit data to one or more servers relating to their motion along road segment 1100 (e.g., steering angle, heading, time, position, speed, sensed road geometry, and/or sensed landmarks, among things). The server may reconstruct trajectories for vehicles 2200 based on the received data. The server may also generate a target trajectory for guiding navigation of an autonomous vehicle that will travel along the same road segment 1100 at a later time based on the first, second, and third trajectories 1101, 1102, and 1103. While a target trajectory may be associated with a single prior traversal of a road segment, in some embodiments, each target trajectory included in sparse map 800 may be determined based on two or more reconstructed trajectories of vehicles traversing the same road segment. In FIG. 11A, the target trajectory is represented by 1110. In some embodiments, the target trajectory 1110 may be generated based on an average of the first, second, and third trajectories 1101, 1102, and 1103. In some embodiments, the target trajectory 1110 included in sparse map 800 may be an aggregation (e.g., a weighted combination) of two or more reconstructed trajectories. Aligning drive data to construct trajectories is further discussed below with respect to FIG. 29.

[0247] FIGS. 11B and 11C further illustrate the concept of target trajectories associated with road segments present within a geographic region 1111. As shown in FIG. 11B, a first road segment 1120 within geographic region 1111 may include a multilane road, which includes two lanes 1122 designated for vehicle travel in a first direction and two additional lanes 1124 designated for vehicle travel in a second direction opposite to the first direction. Lanes 1122 and lanes 1124 may be separated by a double yellow line 1123. Geographic region 1111 may also include a branching road segment 1130 that intersects with road segment 1120. Road segment 1130 may include a two-lane road, each lane being designated for a different direction of travel. Geographic region 1111 may also include other road features, such as a stop line 1132, a stop sign 1134, a speed limit sign 1136, and a hazard sign 1138.

[0248] As shown in FIG. 11C, sparse map 800 may include a local map 1140 including a road model for assisting with autonomous navigation of vehicles within geographic region 1111. For example, local map 1140 may include target trajectories for one or more lanes associated with road segments 1120 and/or 1130 within geographic region 1111. For example, local map 1140 may include target trajectories 1141 and/or 1142 that an autonomous vehicle may access or rely upon when traversing lanes 1122. Similarly, local map 1140 may include target trajectories 1143 and/or 1144 that an autonomous vehicle may access or rely upon when traversing lanes 1124. Further, local map 1140 may include target trajectories 1145 and/or 1146 that an autonomous vehicle may access or rely upon when traversing road segment 1130. Target trajectory 1147 represents a preferred path an autonomous vehicle should follow when transitioning from lanes 1120 (and specifically, relative to target trajectory 1141 associated with a

right-most lane of lanes 1120) to road segment 1130 (and specifically, relative to a target trajectory 1145 associated with a first side of road segment 1130. Similarly, target trajectory 1148 represents a preferred path an autonomous vehicle should follow when transitioning from road segment 1130 (and specifically, relative to target trajectory 1146) to a portion of road segment 1124 (and specifically, as shown, relative to a target trajectory 1143 associated with a left lane of lanes 1124.

[0249] Sparse map 800 may also include representations of other road-related features associated with geographic region 1111. For example, sparse map 800 may also include representations of one or more landmarks identified in geographic region 1111. Such landmarks may include a first landmark 1150 associated with stop line 1132, a second landmark 1152 associated with stop sign 1134, a third landmark 1154 associated with speed limit sign 1154, and a fourth landmark 1156 associated with hazard sign 1138. Such landmarks may be used, for example, to assist an autonomous vehicle in determining its current location relative to any of the shown target trajectories, such that the vehicle may adjust its heading to match a direction of the target trajectory at the determined location. Navigating using landmarks from a sparse map is further discussed below with respect to FIG. 26.

[0250] In some embodiments, sparse map 800 may also include road signature profiles. Such road signature profiles may be associated with any discernible/measurable variation in at least one parameter associated with a road. For example, in some cases, such profiles may be associated with variations in road surface information such as variations in surface roughness of a particular road segment, variations in road width over a particular road segment, variations in distances between dashed lines painted along a particular road segment, variations in road curvature along a particular road segment, etc. FIG. 11D shows an example of a road signature profile 1160. While profile 1160 may represent any of the parameters mentioned above, or others, in one example, profile 1160 may represent a measure of road surface roughness, as obtained, for example, by monitoring one or more sensors providing outputs indicative of an amount of suspension displacement as a vehicle travels a particular road segment.

[0251] Alternatively or concurrently, profile 1160 may represent variation in road width, as determined based on image data obtained via a camera onboard a vehicle traveling a particular road segment. Such profiles may be useful, for example, in determining a particular location of an autonomous vehicle relative to a particular target trajectory. That is, as it traverses a road segment, an autonomous vehicle may measure a profile associated with one or more parameters associated with the road segment. If the measured profile can be correlated/matched with a predetermined profile that plots the parameter variation with respect to position along the road segment, then the measured and predetermined profiles may be used (e.g., by overlaying corresponding sections of the measured and predetermined profiles) in order to determine a current position along the road segment and, therefore, a current position relative to a target trajectory for the road segment.

[0252] In some embodiments, sparse map 800 may include different trajectories based on different characteristics associated with a user of autonomous vehicles, environmental conditions, and/or other parameters relating to driving. For example, in some embodiments, different trajectories may be

generated based on different user preferences and/or profiles. Sparse map 800 including such different trajectories may be provided to different autonomous vehicles or different users. For example, some users may prefer to avoid toll roads, while others may prefer to take the shortest or fastest routes, regardless of whether there is a toll road on the route. The disclosed systems may generate different sparse maps with different trajectories based on such different user preferences or profiles. As another example, some users may prefer to travel in a fast moving lane, while others may prefer to maintain a position in the central lane at all times.

[0253] Different trajectories may be generated and included in sparse map 800 based on different environmental conditions, such as day and night, snow, rain, fog, etc. Autonomous vehicles driving under different environmental conditions may be provided with sparse map 800 generated based on such different environmental conditions. In some embodiments, cameras provided on autonomous vehicles may detect the environmental conditions, and may provide such information back to a server that generates and provides sparse maps. For example, the server may generate or update an already generated sparse map 800 to include trajectories that may be more suitable or safer for autonomous driving under the detected environmental conditions. The update of sparse map 800 based on environmental conditions may be performed dynamically as the autonomous vehicles are traveling along roads.

[0254] Other different parameters relating to driving may also be used as a basis for generating and providing different sparse maps to different autonomous vehicles. For example, when an autonomous vehicle is traveling at a high speed, turns may be tighter. Trajectories associated with specific lanes, rather than roads, may be included in sparse map 800 such that the autonomous vehicle may maintain within a specific lane as the vehicle follows a specific trajectory. When an image captured by a camera onboard the autonomous vehicle indicates that the vehicle has drifted outside of the lane (e.g., crossed the lane mark), an action may be triggered within the vehicle to bring the vehicle back to the designated lane according to the specific trajectory.

[0255] Crowdsourcing a Sparse Map

[0256] In some embodiments, the disclosed systems and methods may generate a sparse for autonomous vehicle navigation. For example, disclosed systems and methods may use crowdsourced data for generation of a sparse, that one or more autonomous vehicles may use to navigate along a system of roads. As used herein, "crowdsourcing" means that data are received from various vehicles (e.g., autonomous vehicles) travelling on a road segment at different times, and such data are used to generate and/or update the road model. The model may, in turn, be transmitted to the vehicles or other vehicles later travelling along the road segment for assisting autonomous vehicle navigation. The road model may include a plurality of target trajectories representing preferred trajectories that autonomous vehicles should follow as they traverse a road segment. The target trajectories may be the same as a reconstructed actual trajectory collected from a vehicle traversing a road segment, which may be transmitted from the vehicle to a server. In some embodiments, the target trajectories may be different from actual trajectories that one or more vehicles previously took when traversing a road segment. The target trajectories may be

generated based on actual trajectories (e.g., through averaging or any other suitable operation). An example of alignment of crowdsourced data for generating target trajectories is discussed below with respect to FIG. 29.

[0257] The vehicle trajectory data that a vehicle may upload to a server may correspond with the actual reconstructed trajectory for the vehicle or may correspond to a recommended trajectory, which may be based on or related to the actual reconstructed trajectory of the vehicle, but may differ from the actual reconstructed trajectory. For example, vehicles may modify their actual reconstructed trajectories and submit (e.g., recommend) to the server the modified actual trajectories. The road model may use the recommended modified trajectories as target trajectories for autonomous navigation of other vehicles.

[0258] In addition to trajectory information, other information for potential use in building a sparse data map 800 may include information relating to potential landmark candidates. For example, through crowdsourcing of information, the disclosed systems and methods may identify potential landmarks in an environment and refine landmark positions. The landmarks may be used by a navigation system of autonomous vehicles to determine and/or adjust the position of the vehicle along the target trajectories.

[0259] The reconstructed trajectories that a vehicle may generate as the vehicle travels along a road may be obtained by any suitable method. In some embodiments, the reconstructed trajectories may be developed by stitching together segments of motion for the vehicle, using, e.g., ego motion estimation (e.g., three dimensional translation and three dimensional rotation of the camera, and hence the body of the vehicle). The rotation and translation estimation may be determined based on analysis of images captured by one or more image capture devices along with information from other sensors or devices, such as inertial sensors and speed sensors. For example, the inertial sensors may include an accelerometer or other suitable sensors configured to measure changes in translation and/or rotation of the vehicle body. The vehicle may include a speed sensor that measures a speed of the vehicle.

[0260] In some embodiments, the ego motion of the camera (and hence the vehicle body) may be estimated based on an optical flow analysis of the captured images. An optical flow analysis of a sequence of images identifies movement of pixels from the sequence of images, and based on the identified movement, determines motions of the vehicle. The ego motion may be integrated over time and along the road segment to reconstruct a trajectory associated with the road segment that the vehicle has followed.

[0261] Data (e.g., reconstructed trajectories) collected by multiple vehicles in multiple drives along a road segment at different times may be used to construct the road model (e.g., including the target trajectories, etc.) included in sparse data map 800. Data collected by multiple vehicles in multiple drives along a road segment at different times may also be averaged to increase an accuracy of the model. In some embodiments, data regarding the road geometry and/or landmarks may be received from multiple vehicles that travel through the common road segment at different times. Such data received from different vehicles may be combined to generate the road model and/or to update the road model.

[0262] The geometry of a reconstructed trajectory (and also a target trajectory) along a road segment may be represented by a curve in three dimensional space, which may be a spline connecting three dimensional polynomials. The reconstructed trajectory curve may be determined from analysis of a video stream or a plurality of images captured by a camera installed on the vehicle. In some embodiments, a location is identified in each frame or image that is a few meters ahead of the current position of the vehicle. This location is where the vehicle is expected to travel to in a predetermined time period. This operation may be repeated frame by frame, and at the same time, the vehicle may compute the camera's ego motion (rotation and translation). At each frame or image, a short range model for the desired path is generated by the vehicle in a reference frame that is attached to the camera. The short range models may be stitched together to obtain a three dimensional model of the road in some coordinate frame, which may be an arbitrary or predetermined coordinate frame. The three dimensional model of the road may then be fitted by a spline, which may include or connect one or more polynomials of suitable orders.

[0263] To conclude the short range road model at each frame, one or more detection modules may be used. For example, a bottom-up lane detection module may be used. The bottom-up lane detection module may be useful when lane marks are drawn on the road. This module may look for edges in the image and assembles them together to form the lane marks. A second module may be used together with the bottom-up lane detection module. The second module is an end-to-end deep neural network, which may be trained to predict the correct short range path from an input image. In both modules, the road model may be detected in the image coordinate frame and transformed to a three dimensional space that may be virtually attached to the camera.

[0264] Although the reconstructed trajectory modeling method may introduce an accumulation of errors due to the integration of ego motion over a long period of time, which may include a noise component, such errors may be inconsequential as the generated model may provide sufficient accuracy for navigation over a local scale. In addition, it is possible to cancel the integrated error by using external sources of information, such as, satellite images or geodetic measurements. For example, the disclosed systems and methods may use a GNSS receiver to cancel accumulated errors. However, the GNSS positioning signals may not be always available and accurate. The disclosed systems and methods may enable a steering application that depends weakly on the availability and accuracy of GNSS positioning. In such systems, the usage of the GNSS signals may be limited. For example, in some embodiments, the disclosed systems may use the GNSS signals for database indexing purposes only.

[0265] In some embodiments, the range scale (e.g., local scale) that may be relevant for an autonomous vehicle navigation steering application may be on the order of 50 meters, 100 meters, 200 meters, 300 meters, etc. Such distances may be used, as the geometrical road model is mainly used for two purposes: planning the trajectory ahead and localizing the vehicle on the road model. In some embodiments, the planning task may use the model over a typical range of 40 meters ahead (or any other suitable distance ahead, such as 20 meters, 30 meters, 50 meters), when the control algorithm steers the vehicle according to a target point located 13 seconds ahead (or any other time such as 1.5 seconds, 1.7

seconds, 2 seconds, etc.). The localization task uses the road model over a typical range of 60 meters behind the car (or any other suitable distances, such as 50 meters, 100 meters, 150 meters, etc.), according to a method called "tail alignment" described in more detail in another section. The disclosed systems and methods may generate a geometrical model that has sufficient accuracy over particular range, such as 100 meters, such that a planned trajectory will not deviate by more than, for example, 330 cm from the lane center.

[0266] As explained above, a three-dimensional road model may be constructed from detecting short range sections and stitching them together. The stitching may be enabled by computing a six degree ego motion model, using the videos and/or images captured by the camera, data from the inertial sensors that reflect the motions of the vehicle, and the host vehicle velocity signal. The accumulated error may be small enough over some local range scale, such as of the order of 100 meters. All this may be completed in a single drive over a particular road segment.

[0267] In some embodiments, multiple drives may be used to average the resulted model, and to increase its accuracy further. The same car may travel the same route multiple times, or multiple cars may send their collected model data to a central server. In any case, a matching procedure may be performed to identify overlapping models and to enable averaging in order to generate target trajectories. The constructed model (e.g., including the target trajectories) may be used for steering once a convergence criterion is met. Subsequent drives may be used for further model improvements and in order to accommodate infrastructure changes.

[0268] Sharing of driving experience (such as sensed data) between multiple cars becomes feasible if they are connected to a central server. Each vehicle client may store a partial copy of a universal road model, which may be relevant for its current position. A bidirectional update procedure between the vehicles and the server may be performed by the vehicles and the server. The small footprint concept discussed above enables the disclosed systems and methods to perform the bidirectional updates using a very small bandwidth.

[0269] Information relating to potential landmarks may also be determined and forwarded to a central server. For example, the disclosed systems and methods may determine one or more physical properties of a potential landmark based on one or more images that include the landmark. The physical properties may include a physical size (e.g., height, width) of the landmark, a distance from a vehicle to a landmark, a distance between the landmark to a previous landmark, the lateral position of the landmark (e.g., the position of the landmark relative to the lane of travel), the GPS coordinates of the landmark, a type of landmark, identification of text on the landmark, etc.. For example, a vehicle may analyze one or more images captured by a camera to detect a potential landmark, such as a speed limit sign.

[0270] The vehicle may determine a distance from the vehicle to the landmark based on the analysis of the one or more images. In some embodiments, the distance may be determined based on analysis of images of the landmark using a suitable image analysis method, such as a scaling method and/or an optical flow method. In some embodiments, the disclosed systems and methods may be configured to determine a type or classification of a potential landmark. In case the vehicle determines

that a certain potential landmark corresponds to a predetermined type or classification stored in a sparse map, it may be sufficient for the vehicle to communicate to the server an indication of the type or classification of the landmark, along with its location. The server may store such indications. At a later time, other vehicles may capture an image of the landmark, process the image (e.g., using a classifier), and compare the result from processing the image to the indication stored in the server with regard to the type of landmark. There may be various types of landmarks, and different types of landmarks may be associated with different types of data to be uploaded to and stored in the server. Different processing onboard the vehicle may detect the landmark and communicate information about the landmark to the server, and the system onboard the vehicle may receive the landmark data from the server and use the landmark data for identifying a landmark in autonomous navigation.

[0271] In some embodiments, multiple autonomous vehicles travelling on a road segment may communicate with a server. The vehicles (or clients) may generate a curve describing its drive (e.g., through ego motion integration) in an arbitrary coordinate frame. The vehicles may detect landmarks and locate them in the same frame. The vehicles may upload the curve and the landmarks to the server. The server may collect data from vehicles over multiple drives, and generate a unified road model. For example, as discussed below with respect to FIG. 19, the server may generate a sparse map having the unified road model using the uploaded curves and landmarks.

[0272] The server may also distribute the model to clients (e.g., vehicles). For example, as discussed below with respect to FIG. 24, the server may distribute the sparse map to one or more vehicles. The server may continuously or periodically update the model when receiving new data from the vehicles. For example, the server may process the new data to evaluate whether the data includes information that should trigger an update, or creation of new data on the server. The server may distribute the updated model or the updates to the vehicles for providing autonomous vehicle navigation.

[0273] The server may use one or more criteria for determining whether new data received from the vehicles should trigger an update to the model or trigger creation of new data. For example, when the new data indicates that a previously recognized landmark at a specific location no longer exists, or is replaced by another landmark, the server may determine that the new data should trigger an update to the model. As another example, when the new data indicates that a road segment has been closed, and when this has been corroborated by data received from other vehicles, the server may determine that the new data should trigger an update to the model.

[0274] The server may distribute the updated model (or the updated portion of the model) to one or more vehicles that are traveling on the road segment, with which the updates to the model are associated. The server may also distribute the updated model to vehicles that are about to travel on the road segment, or vehicles whose planned trip includes the road segment, with which the updates to the model are associated. For example, while an autonomous vehicle is traveling along another road segment before reaching the road segment with which an update is associated, the server may distribute the updates or updated model to the autonomous vehicle before the vehicle reaches the road segment.

5 [(0275)] In some embodiments, the remote server may collect trajectories and landmarks from multiple clients (e.g., vehicles that travel along a common road segment). The server may match curves using landmarks and create an average road model based on the trajectories collected from the multiple vehicles. The server may also compute a graph of roads and the most probable path at each node or conjunction of the road segment. For example, as discussed with respect to FIG. 229 below, the remote server may align the trajectories to generate a crowdsourced sparse map from the collected trajectories.

10 [(0276)] The server may average landmark properties received from multiple vehicles that travelled along the common road segment, such as the distances between one landmark to another (e.g., a previous one along the road segment) as measured by multiple vehicles, to determine an arc-length parameter and support localization along the path and speed calibration for each client vehicle. The server may average the physical dimensions of a landmark measured by multiple vehicles travelled along the common road segment and recognized the same landmark. The averaged physical dimensions may be used to support distance estimation, such as the distance from the vehicle to the landmark. The server may average lateral positions of a landmark (e.g., position from the lane in which vehicles are travelling in to the landmark) measured by multiple vehicles travelled along the common road segment and recognized the same landmark. The averaged lateral position may be used to support lane assignment. The server may average the GPS coordinates of the landmark measured by multiple vehicles travelled along the same road segment and recognized the same landmark. The averaged GPS coordinates of the landmark may be used to support global localization or positioning of the landmark in the road model.

20 [(0277)] In some embodiments, the server may identify model changes, such as constructions, detours, new signs, removal of signs, etc., based on data received from the vehicles. The server may continuously or periodically or instantaneously update the model upon receiving new data from the vehicles. The server may distribute updates to the model or the updated model to vehicles for providing autonomous navigation. For example, as discussed further below, the server may use crowdsourced data to filter out “ghost” landmarks detected by vehicles.

30 [(0278)] In some embodiments, the server may analyze driver interventions during the autonomous driving. The server may analyze data received from the vehicle at the time and location where intervention occurs, and/or data received prior to the time the intervention occurred. The server may identify certain portions of the data that caused or are closely related to the intervention, for example, data indicating a temporary lane closure setup, data indicating a pedestrian in the road. The server may update the model based on the identified data. For example, the server may modify one or more trajectories stored in the model.

35 [(0279)] FIG. 12 is a schematic illustration of a system that uses crowdsourced data to generate a sparse map, as well as distribute and navigate using a crowdsourced sparse map. FIG. 12 shows a road segment 1200 that includes one or more lanes. A plurality of vehicles 1205, 1210, 1215, 1220, and 1225 may travel on road segment 1200 at the same time or at different times (although shown as appearing on road segment 1200 at the same time in FIG. 12). At least one of vehicles 1205, 1210, 1215, 1220, and

1225 may be an autonomous vehicle. For simplicity of the present example, all of the vehicles 1205, 1210, 1215, 1220, and 1225 are presumed to be autonomous vehicles.

[0280] Each vehicle may be similar to vehicles disclosed in other embodiments (e.g., vehicle 200), and may include components or devices included in or associated with vehicles disclosed in other embodiments. Each vehicle may be equipped with an image capture device or camera (e.g., image capture device 122 or camera 122). Each vehicle may communicate with a remote server 1230 via one or more networks (e.g., over a cellular network and/or the Internet, etc.) through wireless communication paths 1235, as indicated by the dashed lines. Each vehicle may transmit data to server 1230 and receive data from server 1230. For example, server 1230 may collect data from multiple vehicles travelling on the road segment 1200 at different times, and may process the collected data to generate an autonomous vehicle road navigation model, or an update to the model. Server 1230 may transmit the autonomous vehicle road navigation model or the update to the model to the vehicles that transmitted data to server 1230. Server 1230 may transmit the autonomous vehicle road navigation model or the update to the model to other vehicles that travel on road segment 1200 at later times.

[0281] As vehicles 1205, 1210, 1215, 1220, and 1225 travel on road segment 1200, navigation information collected (e.g., detected, sensed, or measured) by vehicles 1205, 1210, 1215, 1220, and 1225 may be transmitted to server 1230. In some embodiments, the navigation information may be associated with the common road segment 1200. The navigation information may include a trajectory associated with each of the vehicles 1205, 1210, 1215, 1220, and 1225 as each vehicle travels over road segment 1200. In some embodiments, the trajectory may be reconstructed based on data sensed by various sensors and devices provided on vehicle 1205. For example, the trajectory may be reconstructed based on at least one of accelerometer data, speed data, landmarks data, road geometry or profile data, vehicle positioning data, and ego motion data. In some embodiments, the trajectory may be reconstructed based on data from inertial sensors, such as accelerometer, and the velocity of vehicle 1205 sensed by a speed sensor. In addition, in some embodiments, the trajectory may be determined (e.g., by a processor onboard each of vehicles 1205, 1210, 1215, 1220, and 1225) based on sensed ego motion of the camera, which may indicate three-dimensional translation and/or three-dimensional rotations (or rotational motions). The ego motion of the camera (and hence, the vehicle body) may be determined from analysis of one or more images captured by the camera.

[0282] In some embodiments, the trajectory of vehicle 1205 may be determined by a processor provided aboard vehicle 1205 and transmitted to server 1230. In other embodiments, server 1230 may receive data sensed by the various sensors and devices provided in vehicle 1205, and determine the trajectory based on the data received from vehicle 1205.

[0283] In some embodiments, the navigation information transmitted from vehicles 1205, 1210, 1215, 1220, and 1225 to server 1230 may include data regarding the road surface, the road geometry, or the road profile. The geometry of road segment 1200 may include lane structure and/or landmarks. The lane structure may include the total number of lanes of road segment 1200, the type of lanes (e.g., one-way lane, two-way lane, driving lane, passing lane, etc.), markings on lanes, width of lanes, etc. In some

embodiments, the navigation information may include a lane assignment, e.g., which lane of a plurality of lanes a vehicle is traveling in. For example, the lane assignment may be associated with a numerical value, "3", indicating that the vehicle is traveling on the third lane from the left or right. As another example, the lane assignment may be associated with a text value "center lane" indicating the vehicle is traveling on the center lane.

[0284] Server 1230 may store the navigation information on a non-transitory computer-readable medium, such as a hard drive, a compact disc, a tape, a memory, etc. Server 1230 may generate (e.g., through a processor included in server 1230) at least a portion of an autonomous vehicle road navigation model for the common road segment 1200 based on the navigation information received from the plurality of vehicles 1205, 1210, 1215, 1220, and 1225 and may store the model as a portion of a sparse map. Server 1230 may determine a trajectory associated with each lane based on crowdsourced data (e.g., navigation information) received from multiple vehicles (e.g., 1205, 1210, 1215, 1220, and 1225) that travel on a lane of road segment at different times. Server 1230 may generate the autonomous vehicle road navigation model or a portion of the model (e.g., an updated portion) based on a plurality of trajectories determined based on the crowd sourced navigation data. As explained in greater detail below with respect to FIG. 24, server 1230 may transmit the model or the updated portion of the model to one or more of autonomous vehicles 1205, 1210, 1215, 1220, and 1225 traveling on road segment 1200 or any other autonomous vehicles that travel on road segment at a later time for updating an existing autonomous vehicle road navigation model provided in a navigation system of the vehicles. As explained in greater detail below with respect to FIG. 26, the autonomous vehicle road navigation model may be used by the autonomous vehicles in autonomously navigating along the common road segment 1200.

[0285] As explained above, the autonomous vehicle road navigation model may be included in a sparse map (e.g., sparse map 800 depicted in FIG. 8). Sparse map 800 may include sparse recording of data related to road geometry and/or landmarks along a road, which may provide sufficient information for guiding autonomous navigation of an autonomous vehicle, yet does not require excessive data storage. In some embodiments, the autonomous vehicle road navigation model may be stored separately from sparse map 800, and may use map data from sparse map 800 when the model is executed for navigation. In some embodiments, the autonomous vehicle road navigation model may use map data included in sparse map 800 for determining target trajectories along road segment 1200 for guiding autonomous navigation of autonomous vehicles 1205, 1210, 1215, 1220, and 1225 or other vehicles that later travel along road segment 1200. For example, when the autonomous vehicle road navigation model is executed by a processor included in a navigation system of vehicle 1205, the model may cause the processor to compare the trajectories determined based on the navigation information received from vehicle 1205 with predetermined trajectories included in sparse map 800 to validate and/or correct the current traveling course of vehicle 1205.

[0286] In the autonomous vehicle road navigation model, the geometry of a road feature or target trajectory may be encoded by a curve in a three-dimensional space. In one embodiment, the curve may be a three dimensional spline including one or more connecting three dimensional polynomials. As

one of skill in the art, would understand, a spline may be a numerical function that is piece-wise defined by a series of polynomials for fitting data. A spline for fitting the three dimensional geometry data of the road may include a linear spline (first order), a quadratic spline (second order), a cubic spline (third order), or any other splines (other orders), or a combination thereof. The spline may include one or more three dimensional polynomials of different orders connecting (e.g., fitting) data points of the three dimensional geometry data of the road. In some embodiments, the autonomous vehicle road navigation model may include a three dimensional spline corresponding to a target trajectory along a common road segment (e.g., road segment 1200) or a lane of the road segment 1200.

[0287] As explained above, the autonomous vehicle road navigation model included in the sparse map may include other information, such as identification of at least one landmark along road segment 1200. The landmark may be visible within a field of view of a camera (e.g., camera 1122) installed on each of vehicles 1205, 1210, 1215, 1220, and 1225. In some embodiments, camera 1122 may capture an image of a landmark. A processor (e.g., processor 1180, 1190, or processing unit 110) provided on vehicle 1205 may process the image of the landmark to extract identification information for the landmark. The landmark identification information, rather than an actual image of the landmark, may be stored in sparse map 800. The landmark identification information may require much less storage space than an actual image. Other sensors or systems (e.g., GPS system) may also provide certain identification information of the landmark (e.g., position of landmark). The landmark may include at least one of a traffic sign, an arrow marking, a lane marking, a dashed lane marking, a traffic light, a stop line, a directional sign (e.g., a highway exit sign with an arrow indicating a direction, a highway sign with arrows pointing to different directions or places), a landmark beacon, or a lamppost. A landmark beacon refers to a device (e.g., an RFID device) installed along a road segment that transmits or reflects a signal to a receiver installed on a vehicle, such that when the vehicle passes by the device, the beacon received by the vehicle and the location of the device (e.g., determined from GPS location of the device) may be used as a landmark to be included in the autonomous vehicle road navigation model and/or the sparse map 800.

[0288] The identification of at least one landmark may include a position of the at least one landmark. The position of the landmark may be determined based on position measurements performed using sensor systems (e.g., Global Positioning Systems, inertial based positioning systems, landmark beacon, etc.) associated with the plurality of vehicles 1205, 1210, 1215, 1220, and 1225. In some embodiments, the position of the landmark may be determined by averaging the position measurements detected, collected, or received by sensor systems on different vehicles 1205, 1210, 1215, 1220, and 1225 through multiple drives. For example, vehicles 1205, 1210, 1215, 1220, and 1225 may transmit position measurements data to server 1230, which may average the position measurements and use the averaged position measurement as the position of the landmark. The position of the landmark may be continuously refined by measurements received from vehicles in subsequent drives.

[0289] The identification of the landmark may include a size of the landmark. The processor provided on a vehicle (e.g., 1205) may estimate the physical size of the landmark based on the analysis of

the images. Server 1230 may receive multiple estimates of the physical size of the same landmark from different vehicles over different drives. Server 1230 may average the different estimates to arrive at a physical size for the landmark, and store that landmark size in the road model. The physical size estimate may be used to further determine or estimate a distance from the vehicle to the landmark. The distance to the landmark may be estimated based on the current speed of the vehicle and a scale of expansion based on the position of the landmark appearing in the images relative to the focus of expansion of the camera. For example, the distance to landmark may be estimated by $Z = V * dt * R / D$, where V is the speed of vehicle, R is the distance in the image from the landmark at time t_1 to the focus of expansion, and D is the change in distance for the landmark in the image from t_1 to t_2 . dt represents the $(t_2 - t_1)$. For example, the distance to landmark may be estimated by $Z = V * dt * R / D$, where V is the speed of vehicle, R is the distance in the image between the landmark and the focus of expansion, dt is a time interval, and D is the image displacement of the landmark along the epipolar line. Other equations equivalent to the above equation, such as $Z = V * \omega / \Delta \omega$, may be used for estimating the distance to the landmark. Here, V is the vehicle speed, ω is an image length (like the object width), and $\Delta \omega$ is the change of that image length in a unit of time.

[0290] When the physical size of the landmark is known, the distance to the landmark may also be determined based on the following equation: $Z = f * W / \omega$, where f is the focal length, W is the size of the landmark (e.g., height or width), ω is the number of pixels when the landmark leaves the image. From the above equation, a change in distance Z may be calculated using $\Delta Z \equiv f * \ddot{W} * \Delta \omega / \omega^2 + f * \Delta W / \omega$, where ΔW decays to zero by averaging, and where $\Delta \omega$ is the number of pixels representing a bounding box accuracy in the image. A value estimating the physical size of the landmark may be calculated by averaging multiple observations at the server side. The resulting error in distance estimation may be very small. There are two sources of error that may occur when using the formula above, namely ΔW and $\Delta \omega$. Their contribution to the distance error is given by $\Delta Z \equiv f * \ddot{W} * \Delta \omega / \omega^2 + f * \Delta W / \omega$. However, ΔW decays to zero by averaging; hence ΔZ is determined by $\Delta \omega$ (e.g., the inaccuracy of the bounding box in the image).

[0291] For landmarks of unknown dimensions, the distance to the landmark may be estimated by tracking feature points on the landmark between successive frames. For example, certain features appearing on a speed limit sign may be tracked between two or more image frames. Based on these tracked features, a distance distribution per feature point may be generated. The distance estimate may be extracted from the distance distribution. For example, the most frequent distance appearing in the distance distribution may be used as the distance estimate. As another example, the average of the distance distribution may be used as the distance estimate.

[0292] FIG. 13 illustrates an example autonomous vehicle road navigation model represented by a plurality of three-dimensional splines 1301, 1302, and 1303. The curves 1301, 1302, and 1303 shown in FIG. 13 are for illustration purpose only. Each spline may include one or more three-dimensional polynomials connecting a plurality of data points 1310. Each polynomial may be a first order polynomial, a second order polynomial, a third order polynomial, or a combination of any suitable

polynomials having different orders. Each data point 1310 may be associated with the navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225. In some embodiments, each data point 1310 may be associated with data related to landmarks (e.g., size, location, and identification information of landmarks) and/or road signature profiles (e.g., road geometry, road roughness profile, road curvature profile, road width profile). In some embodiments, some data points 1310 may be associated with data related to landmarks, and others may be associated with data related to road signature profiles.

[0293] FIG. 14 illustrates raw location data 1410 (e.g., GPS data) received from five separate drives. One drive may be separate from another drive if it was traversed by separate vehicles at the same time, by the same vehicle at separate times, or by separate vehicles at separate times. To account for errors in the location data 1410 and for differing locations of vehicles within the same lane (e.g., one vehicle may drive closer to the left of a lane than another), server 1230 may generate a map skeleton 1420 using one or more statistical techniques to determine whether variations in the raw location data 1410 represent actual divergences or statistical errors. Each path within skeleton 1420 may be linked back to the raw data 1410 that formed the path. For example, the path between A and B within skeleton 1420 is linked to raw data 1410 from drives 2, 3, 4, and 5 but not from drive 1. Skeleton 1420 may not be detailed enough to be used to navigate a vehicle (e.g., because it combines drives from multiple lanes on the same road unlike the splines described above) but may provide useful topological information and may be used to define intersections.

[0294] FIG. 15 illustrates an example by which additional detail may be generated for a sparse map within a segment of a map skeleton (e.g., segment A to B within skeleton 1420). As depicted in FIG. 15, the data (e.g., ego-motion data, road markings data, and the like) may be shown as a function of position S (or S_i or A_i) along the drive. Server 1230 may identify landmarks for the sparse map by identifying unique matches between landmarks 1501, 1503, and 1505 of drive 1510 and landmarks 1507 and 1509 of drive 1520. Such a matching algorithm may result in identification of landmarks 1511, 1513, and 1515. One skilled in the art would recognize, however, that other matching algorithms may be used. For example, probability optimization may be used in lieu of or in combination with unique matching. As described in further detail below with respect to FIG. 29, server 1230 may longitudinally align the drives to align the matched landmarks. For example, server 1230 may select one drive (e.g., drive 1520) as a reference drive and then shift and/or elastically stretch the other drive(s) (e.g., drive 1510) for alignment.

[0295] FIG. 16 shows an example of aligned landmark data for use in a sparse map. In the example of FIG. 16, landmark 1610 comprises a road sign. The example of FIG. 16 further depicts data from a plurality of drives 1601, 1603, 1605, 1607, 1609, 1611 and 1613. In the example of FIG. 16, the data from drive 1613 consists of a "ghost" landmark, and the server 1230 may identify it as such because none of drives 1601, 1603, 1605, 1607, 1609, and 1611 include an identification of a landmark in the vicinity of the identified landmark in drive 1613. Accordingly, server 1230 may accept potential landmarks when a ratio of images in which the landmark does appear to images in which the landmark

does not appear exceeds a threshold and/or may reject potential landmarks when a ratio of images in which the landmark does not appear to images in which the landmark does appear exceeds a threshold.

[0296] FIG. 17 depicts a system 1700 for generating drive data, which may be used to crowdsource a sparse map. As depicted in FIG. 17, system 1700 may include a camera 1701a and a locating device 1703 (e.g., a GPS locator). Camera 1701a and locating device 1703 may be mounted on a vehicle (e.g., one of vehicles 1205, 1210, 1215, 1220, and 1225). Camera 1701a may produce a plurality of data of multiple types, e.g., ego motion data, traffic sign data, road data, or the like. The camera data and location data may be segmented into drive segments 1705. For example, drive segments 1705 may each have camera data and location data from less than 1 km of driving.

[0297] In some embodiments, system 1700 may remove redundancies in drive segments 1705. For example, if a landmark appears in multiple images from camera 1701, system 1700 may strip the redundant data such that the drive segments 1705 only contain one copy of the location of and any metadata relating to the landmark. By way of further example, if a lane marking appears in multiple images from camera 1701, system 1700 may strip the redundant data such that the drive segments 1705 only contain one copy of the location of and any metadata relating to the lane marking.

[0298] System 1700 also includes a server (e.g., server 1230). Server 1230 may receive drive segments 1705 from the vehicle and recombine the drive segments 1705 into a single drive 1707. Such an arrangement may allow for reduce bandwidth requirements when transferring data between the vehicle and the server while also allowing for the server to store data relating to an entire drive.

[0299] FIG. 18 depicts system 1700 of FIG. 17 further configured for crowdsourcing a sparse map. As in FIG. 17, system 1700 includes vehicle 1810, which captures drive data using, for example, a camera (which produces, e.g., ego motion data, traffic sign data, road data, or the like) and a locating device (e.g., a GPS locator). As in FIG. 17, vehicle 1810 segments the collected data into drive segments (depicted as “DS1 1,” “DS2 1,” “DSN 1” in FIG. 18). Server 1230 then receives the drive segments and reconstructs a drive (depicted as “Drive 1” in FIG. 18) from the received segments.

[0300] As further depicted in FIG. 18, system 1700 also receives data from additional vehicles. For example, vehicle 1820 also captures drive data using, for example, a camera (which produces, e.g., ego motion data, traffic sign data, road data, or the like) and a locating device (e.g., a GPS locator). Similar to vehicle 1810, vehicle 1820 segments the collected data into drive segments (depicted as “DS1 2,” “DS2 2,” “DSN 2” in FIG. 18). Server 1230 then receives the drive segments and reconstructs a drive (depicted as “Drive 2” in FIG. 18) from the received segments. Any number of additional vehicles may be used. For example, FIG. 18 also includes “CAR N” that captures drive data segments into drive segments (depicted as “DS1 N,” “DS2 N,” “DSN N” in FIG. 18) and sends it to server 1230 for reconstruction into a drive (depicted as “Drive N” in FIG. 18).

[0301] As depicted in FIG. 18, server 1230 may construct a sparse map (depicted as “MAP”) using the reconstructed drives (e.g., “Drive 1,” “Drive 2,” and “Drive N”) collected from a plurality of vehicles (e.g., “CAR 1” (also labeled vehicle 1810), “CAR 2” (also labeled vehicle 1820), and “CAR N”).

[0302] FIG. 19 is a flowchart showing an example process 1900 for generating a sparse map for autonomous vehicle navigation along a road segment. Process 1900 may be performed by one or more processing devices included in server 1230.

[0303] Process 1900 may include receiving a plurality of images acquired as one or more vehicles traverse the road segment (step 1905). Server 1230 may receive images from cameras included within one or more of vehicles 1205, 1210, 1215, 1220, and 1225. For example, camera 122 may capture one or more images of the environment surrounding vehicle 1205 as vehicle 1205 travels along road segment 1200. In some embodiments, server 1230 may also receive stripped down image data that has had redundancies removed by a processor on vehicle 1205, as discussed above with respect to FIG. 17.

[0304] Process 1900 may further include identifying, based on the plurality of images, at least one line representation of a road surface feature extending along the road segment (step 1910). Each line representation may represent a path along the road segment substantially corresponding with the road surface feature. For example, server 1230 may analyze the environmental images received from camera 122 to identify a road edge or a lane marking and determine a trajectory of travel along road segment 1200 associated with the road edge or lane marking. In some embodiments, the trajectory (or line representation) may include a spline, a polynomial representation, or a curve. Server 1230 may determine the trajectory of travel of vehicle 1205 based on camera ego motions (e.g., three dimensional translation and/or three dimensional rotational motions) received at step 1905.

[0305] Process 1900 may also include identifying, based on the plurality of images, a plurality of landmarks associated with the road segment (step 1910). For example, server 1230 may analyze the environmental images received from camera 122 to identify one or more landmarks, such as road sign along road segment 1200. Server 1230 may identify the landmarks using analysis of the plurality of images acquired as one or more vehicles traverse the road segment. To enable crowdsourcing, the analysis may include rules regarding accepting and rejecting possible landmarks associated with the road segment. For example, the analysis may include accepting potential landmarks when a ratio of images in which the landmark does appear to images in which the landmark does not appear exceeds a threshold and/or rejecting potential landmarks when a ratio of images in which the landmark does not appear to images in which the landmark does appear exceeds a threshold.

[0306] Process 1900 may include other operations or steps performed by server 1230. For example, the navigation information may include a target trajectory for vehicles to travel along a road segment, and process 1900 may include clustering, by server 1230, vehicle trajectories related to multiple vehicles travelling on the road segment and determining the target trajectory based on the clustered vehicle trajectories, as discussed in further detail below. Clustering vehicle trajectories may include clustering, by server 1230, the multiple trajectories related to the vehicles travelling on the road segment into a plurality of clusters based on at least one of the absolute heading of vehicles or lane assignment of the vehicles. Generating the target trajectory may include averaging, by server 1230, the clustered trajectories.

[(0307)] By way of further example, process 1900 may include aligning data received in step 1905, as discussed in further detail below with respect to FIG. 29. Other processes or steps performed by server 1230, as described above, may also be included in process 1900.

[(0308)] The disclosed systems and methods may include other features. For example, the disclosed systems may use local coordinates, rather than global coordinates. For autonomous driving, some systems may present data in world coordinates. For example, longitude and latitude coordinates on the earth's surface may be used. In order to use the map for steering, the host vehicle may determine its position and orientation relative to the map. It seems natural to use a GPS device on board, in order to position the vehicle on the map and in order to find the rotation transformation between the body reference frame and the world reference frame (e.g., North, East and Down). Once the body reference frame is aligned with the map reference frame, then the desired route may be expressed in the body reference frame and the steering commands may be computed or generated.

[(0309)] However, one possible issue with this strategy is that current GPS technology does not usually provide the body location and position with sufficient accuracy and availability. To overcome this problem, landmarks whose world coordinates are known may be used to construct very detailed maps (called High Definition or HD maps), that contain landmarks of different kinds. Accordingly, a vehicle equipped with a sensor may detect and locate the landmarks in its own reference frame. Once the relative position between the vehicle and the landmarks is found, the landmarks' world coordinates may be determined from the HD map, and the vehicle may use them to compute its own location and position.

[(0310)] This method may nevertheless use the global world coordinate system as a mediator that establishes the alignment between the map and the body reference frames. Namely, the landmarks may be used in order to compensate for the limitations of the GPS device onboard the vehicles. The landmarks, together with an HD map, may enable to compute the precise vehicle position in global coordinates, and hence the map-body alignment problem is solved.

[(0311)] In the disclosed systems and methods, instead of using one global map of the world, many map pieces or local maps may be used for autonomous navigation. Each piece of a map or each local map may define its own coordinate frame. These coordinate frames may be arbitrary. The vehicle's coordinates in the local maps may not need to indicate where the vehicle is located on the surface of earth. Moreover, the local maps may not be required to be accurate over large scales, meaning there may be no rigid transformation that can embed a local map in the global world coordinate system.

[(0312)] There are two main processes associated with this representation of the world, one relates to the generation of the maps and the other relates to using them. With respect to map generation, this type of representation may be created and maintained by crowdsourcing. There may be no need to apply sophisticated survey equipment, because the use of HD maps is limited, and hence crowdsourcing becomes feasible. With respect to usage, an efficient method to align the local map with the body reference frame without going through a standard world coordinate system may be employed. Hence there may be no need, at least in most scenarios and circumstances, to have a precise estimation of the

vehicle location and position in global coordinates. Further, the memory footprint of the local maps may be kept very small.

[0313] The principle underlying the maps generation is the integration of ego motion. The vehicles may sense the motion of the camera in space (3D translation and 3D rotation). The vehicles or the server may reconstruct the trajectory of the vehicle by integration of ego motion over time, and this integrated path may be used as a model for the road geometry. This process may be combined with sensing of close-range lane marks, and then the reconstructed route may reflect the path that a vehicle should follow, and not the particular path that the vehicle did follow. In other words, the reconstructed route or trajectory may be modified based on the sensed data relating to close range lane marks, and the modified reconstructed trajectory may be used as a recommended trajectory or target trajectory, which may be saved in the road model or sparse map for use by other vehicles navigating the same road segment.

[0314] In some embodiments, the map coordinate system may be arbitrary. A camera reference frame may be selected at an arbitrary time, and used as the map origin. The integrated trajectory of the camera may be expressed in the coordinate system of that particular chosen frame. The value of the route coordinates in the map may not directly represent a location on earth.

[0315] The integrated path may accumulate errors. This may be due to the fact that the sensing of the ego motion may not be absolutely accurate. The result of the accumulated error is that the local map may diverge, and the local map may not be regarded as a local copy of the global map. The larger the size of the local map piece, the larger the deviation from the "true" geometry on earth.

[0316] The arbitrariness and the divergence of the local maps may not be a consequence of the integration method, which may be applied in order to construct the maps in a crowdsourcing manner (e.g., by vehicles traveling along the roads). However, vehicles may successfully use the local maps for steering.

[0317] The map may diverge over long distances. Since the map is used to plan a trajectory in the immediate vicinity of the vehicle, the effect of the divergence may be acceptable. At any time instance, the system (e.g., server 1230 or vehicle 1205) may repeat the alignment procedure, and use the map to predict the road location (in the camera coordinate frame) some 1.3 seconds ahead (or any other seconds, such as 1.5 seconds, 1.0 second, 1.8 seconds, etc.). As long as the accumulated error over that distance is small enough, then the steering command provided for autonomous driving may be used.

[0318] In some embodiments, a local map may focus on a local area, and may not cover a too large area. This means that a vehicle that is using a local map for steering in autonomous driving, may arrive at some point to the end of the map and may have to switch to another local piece or section of a map. The switching may be enabled by the local maps overlapping each other. Once the vehicle enters the area that is common to both maps, the system (e.g., server 1230 or vehicle 1205) may continue to generate steering commands based on a first local map (the map that is being used), but at the same time the system may localize the vehicle on the other map (or second local map) that overlaps with the first local map. In other words, the system may simultaneously align the present coordinate frame of the

camera both with the coordinate frame of the first map and with the coordinate frame of the second map. When the new alignment is established, the system may switch to the other map and plan the vehicle trajectory there.

[0319] The disclosed systems may include additional features, one of which is related to the way the system aligns the coordinate frames of the vehicle and the map. As explained above, landmarks may be used for alignment, assuming the vehicle may measure its relative position to them. This is useful in autonomous driving, but sometimes it may result in a demand for a large number of landmarks and hence a large memory footprint. The disclosed systems may therefore use an alignment procedure that addresses this problem. In the alignment procedure, the system may compute a 1D estimator for the location of the vehicle along the road, using sparse landmarks and integration of ego speed. The system may use the shape of the trajectory itself to compute the rotation part of the alignment, using a tail alignment method discussed in details below in other sections. Accordingly, the vehicle may reconstruct its own trajectory while driving the “tail” and computes a rotation around its assumed position along the road, in order to align the tail with the map. Such an alignment procedure is distinct from the alignment of the crowdsourced data discussed below with respect to FIG. 29.

[0320] In the disclosed systems and methods, a GPS device may still be used. Global coordinates may be used for indexing the database that stores the trajectories and/or landmarks. The relevant piece of local map and the relevant landmarks in the vicinity of the vehicles may be stored in memory and retrieved from the memory using global GPS coordinates. However, in some embodiments, the global coordinates may not be used for path planning, and may not be accurate. In one example, the usage of global coordinates may be limited for indexing of the information.

[0321] In situations where “tail alignment” cannot function well, the system may compute the vehicle’s position using a larger number of landmarks. This may be a rare case, and hence the impact on the memory footprint may be moderate. Road intersections are examples of such situations.

[0322] The disclosed systems and methods may use semantic landmarks (e.g., traffic signs), since they can be reliably detected from the scene and matched with the landmarks stored in the road model or sparse map. In some cases, the disclosed systems may use non-semantic landmarks (e.g., general purpose signs) as well, and in such cases the non-semantic landmarks may be attached to an appearance signature, as discussed above. The system may use a learning method for the generation of signatures that follows the “same-or-not-same” recognition paradigm.

[0323] For example, given many drives with GPS coordinates along them, the disclosed systems may produce the underlying road structure: junctions and road segments, as discussed above with respect to FIG. 14. The roads are assumed to be far enough from each other to be able to differentiate them using the GPS. In some embodiments, only a coarse grained map may be needed. To generate the underlying road structure graph, the space may be divided into a lattice of a given resolution (e.g., 50m by 50m). Every drive may be seen as an ordered list of lattice sites. The system may color every lattice site belonging to a drive to produce an image of the merged drives. The colored lattice points may be represented as nodes on the merged drives. The drives passing from one node to another may be

represented as links. The system may fill small holes in the image, to avoid differentiating lanes and correct for GPS errors. The system may use a suitable thinning algorithm (e.g., an algorithm named “Zhang-Suen” thinning algorithm) to obtain the skeleton of the image. This skeleton may represent the underlying road structure, and junctions may be found using a mask (e.g., a point connected to at least three others). After the junctions are found, the segments may be the skeleton parts that connect them. To match the drives back to the skeleton, the system may use a Hidden Markov Model. Every GPS point may be associated with a lattice site, with a probability inverse to its distance from that site. Use a suitable algorithm (e.g., an algorithm named the “Viterbi” algorithm) to match GPS points to lattice sites, while not allowing consecutive GPS points to match to non-neighboring lattice sites.

[0324] A plurality of methods may be used for mapping the drives back to the map. For example, a first solution may include keeping track during the thinning process. A second solution may use proximity matching. A third solution may use hidden Markov model. The hidden Markov model assumes an underlying hidden state for every observation, and assigns probabilities for a given observation given the state, and for a state given the previous state. A Viterbi algorithm may be used to find the most probable states given a list of observations.

[0325] The disclosed systems and methods may include additional features. For example, the disclosed systems and methods may detect highway entrances/exits. Multiple drives in the same area may be merged using GPS data to the same coordinate system. The system may use visual feature points for mapping and localization.

[0326] In some embodiments, generic visual features may be used as landmarks for the purpose of registering the position and orientation of a moving vehicle, in one drive (localization phase), relative to a map generated by vehicles traversing the same stretch of road in previous drives (mapping phase). These vehicles may be equipped with calibrated cameras imaging the vehicle surroundings and GPS receivers. The vehicles may communicate with a central server (e.g., server 1230) that maintains an up-to-date map including these visual landmarks, connected to other significant geometric and semantic information (e.g., lane structure, type, and position of road signs, type and position of road marks, shape of nearby drivable ground area delineated by the position of physical obstacles, shape of previously driven vehicle path when controlled by human driver, etc.). The total amount of data that may be communicated between the central server and vehicles per length of road is small, both in a mapping and localization phases.

[0327] In a mapping phase, disclosed systems (e.g., autonomous vehicles and/or one or more servers) may detect feature points (FPs). Feature points may comprise one or more points that are used to track an associated object such as a landmark. For example, the eight points comprising the corners of a stop sign may be feature points. Disclosed systems may further compute descriptors associated with the FPs (e.g., using the features from the accelerated segment test (FAST) detector, the binary robust invariant scalable keypoints (BRISK) detector, the binary robust independent elementary features (BRIEF) detector, and/or the oriented FAST and rotated BRIEF (ORB) detector or using a detector/descriptor pair that was trained using a training library). The system may track FPs between

frames in which they appear using their motion in the image plane and by matching the associated descriptors using, for example, Euclidean or Hamming distance in descriptor space. The system may use tracked FPs to estimate camera motion and world positions of objects on which FPs were detected and tracked. For example, tracked FPs may be used to estimate the motion of the vehicle and/or the position of a landmark on which the FPs were initially detected.

[0328] The system may further classify FPs as ones that will likely be detected in future drives or not (e.g., FPs detected on momentarily moving objects, parked cars, and shadow texture will likely not reappear in future drives). This classification may be referred to as a reproducibility classification (RC) and may be a function of the intensities of the light in a region of a pyramid surrounding the detected FP, the motion of the tracked FP in the image plane, and/or the extent of viewpoints in which it was successfully detected and tracked. In some embodiments, the vehicles may send descriptors associated with an FP, estimated 3D position relative to the vehicle of the FP, and momentary vehicle GPS coordinates at the time of detecting/tracking the FP, to server 11230.

[0329] During a mapping phase, when communication bandwidth between mapping vehicles and a central server is limited, the vehicles may send FPs to the server at a high frequency when the presence of FPs or other semantic landmarks in the map (such as road signs and lane structure) is limited and insufficient for the purpose of localization. Moreover, although vehicles in the mapping phase may generally send FPs to the server at a low spatial frequency, the FPs may be agglomerated in the server. Detection of reoccurring FPs may also be performed by the server and the server may store the set of reoccurring FPs and/or disregard FPs that do not reoccur. Visual appearance of landmarks may, at least in some cases, be sensitive to the time of day or the season in which they were captured. Accordingly, to increase reproducibility probability of FPs, the received FPs may be binned by the server into time-of-day bins, season bins, and the like. In some embodiments, the vehicles may also send the server other semantic and geometric information associated with the FPs (e.g., lane shape, structure of road plane, 3D position of obstacles, free space in mapping clip, momentary coordinate system, path driven by human driver in a setup, drive to a parking location, etc.).

[0330] In a localization phase, the server may send a map containing landmarks in the form of FP positions and descriptors to one or more vehicles. Feature points (FPs) may be detected and tracked by the vehicles in near-real time within a set of current consecutive frames. Tracked FPs may be used to estimate camera motion and/or positions of associated objects such as landmarks. Detected FP descriptors may be searched to match a list of FPs included in the map and having GPS coordinates within an estimated finite GPS uncertainty radius from the current GPS reading of the vehicle. Matching may be done by searching all pairs of current and mapping FPs that minimize an Euclidean or Hamming distance in descriptor space. Using the FP matches and their current and map positions, the vehicle may rotate and/or translate between the momentary vehicle position and the local map coordinate system.

[0331] The disclosed systems and methods may include a method for training a reproducibility classifier. Training may be performed in one of the following schemes in order to grow a labeling cost and resulting classifier accuracy.

[0332] In the first scheme, a database including a large number of clips recorded by vehicle cameras with matching momentary vehicle GPS position may be collected. This database may include a representative sample of drives (with respect to various properties: e.g., time of day, season, weather condition, type of roadway). Feature points (FPs) extracted from frames of different drives at a similar GPS position and heading may be potentially matched within a GPS uncertainty radius. Unmatched FPs may be labeled unreproducible and those matched may be labeled reproducible. A classifier may then be trained to predict the reproducibility label of an FP given its appearance in the image pyramid, its momentary position relative to the vehicle and the extent of viewpoints positions in which it was successfully tracked.

[0333] In the second scheme, FP pairs extracted from the clip database described in the first scheme may also be labeled by a human responsible for annotating FP matches between clips.

[0334] In a third scheme, a database augmenting that of the first scheme with precise vehicle position, vehicle orientation and image pixel depth using Light Detection And Ranging (LIDAR) measurements may be used to accurately match world positions in different drives. Feature point descriptors may then be computed at the image region corresponding to these world points at different viewpoints and drive times. The classifier may then be trained to predict the average distance in descriptor space a descriptor is located from its matched descriptors. In this case reproducibility may be measured by likely having a low descriptor distance.

[0335] Consistent with disclosed embodiments, the system may generate an autonomous vehicle road navigation model based on the observed trajectories of vehicles traversing a common road segment (e.g., which may correspond to the trajectory information forwarded to a server by a vehicle). The observed trajectories, however, may not correspond to actual trajectories taken by vehicles traversing a road segment. Rather, in certain situations, the trajectories uploaded to the server may be modified with respect to actual reconstructed trajectories determined by the vehicles. For example, a vehicle system, while reconstructing a trajectory actually taken, may use sensor information (e.g., analysis of images provided by a camera) to determine that its own trajectory may not be the preferred trajectory for a road segment. For example, the vehicle may determine based on image data from onboard cameras that it is not driving in a center of a lane or that it crossed over a lane boundary for a determined period of time. In such cases, among others, a refinement to the vehicle's reconstructed trajectory (the actual path traversed) may be made based on information derived from the sensor output. The refined trajectory, not the actual trajectory, may then be uploaded to the server for potential use in building or updating a sparse data map.

[0336] In some embodiments, then, a processor included in a vehicle (e.g., vehicle 1205) may determine an actual trajectory of vehicle 1205 based on the outputs from the one or more sensors. For example, based on analysis of images output from camera 122, the processor may identify landmarks along road segment 1200. Landmarks may include traffic signs (e.g., speed limit signs), directional signs (e.g., highway directional signs pointing to different routes or places), and general signs (e.g., a rectangular business sign that is associated with a unique signature, such as a color pattern). The

identified landmark may be compared with the landmark stored in sparse map 800. When a match is found, the location of the landmark stored in sparse map 800 may be used as the location of the identified landmark. The location of the identified landmark may be used for determining the location of the vehicle 1205 along a target trajectory. In some embodiments, the processor may also determine the location of vehicle 1205 based on GPS signals output by GPS unit 1710.

[0337] The processor may also determine a target trajectory for transmitting to server 1230. The target trajectory may be the same as the actual trajectory determined by the processor based on the sensor outputs. In some embodiments, though, the target trajectory may be different from the actual trajectory determined based on the sensor outputs. For example, the target trajectory may include one or more modifications to the actual trajectory.

[0338] In one example, if data from camera 1122 includes a barrier, such as a temporary lane shifting barrier 100 meters ahead of vehicle 11250 that changes the lanes (e.g., when lanes are temporarily shifted due to constructions or an accident ahead), the processor may detect the temporary lane shifting barrier from the image, and select a lane different from a lane corresponding to the target trajectory stored in the road model or sparse map in compliance to the temporary lane shift. The actual trajectory of vehicle 1205 may reflect this change of lanes. However, if the lane shifting is temporary and may be cleared in the next 10, 15, or 30 minutes, for example, vehicle 1205 may thus modify the actual trajectory (i.e., the shift of lanes) vehicle 1205 has taken to reflect that a target trajectory should be different from the actual trajectory vehicle 1205 has taken. For example, the system may recognize that the path traveled differs from a preferred trajectory for the road segment. Thus, the system may adjust a reconstructed trajectory prior to uploading the trajectory information to the servers.

[0339] In other embodiments, the actual reconstructed trajectory information may be uploaded, by one or more recommended trajectory refinements (e.g., a size and direction of a translation to be made to at least a portion of the reconstructed trajectory) may also be uploaded. In some embodiments, processor 1715 may transmit a modified actual trajectory to server 1230. Server 1230 may generate or update a target trajectory based on the received information and may transmit the target trajectory to other autonomous vehicles that later travel on the same road segment, as discussed in further detail below with respect to FIG. 24.

[0340] As another example, the environmental image may include an object, such as a pedestrian suddenly appearing in road segment 1200. The processor may detect the pedestrian, and vehicle 1205 may change lanes to avoid a collision with the pedestrian. The actual trajectory of vehicle 1205 reconstructed based on sensed data may include the change of lanes. However, the pedestrian may soon leave the roadway. So, vehicle 1205 may modify the actual trajectory (or determine a recommended modification) to reflect that the target trajectory should be different from the actual trajectory taken (as the appearance of the pedestrian is a temporary condition that should not be accounted for in the target trajectory determination). In some embodiments, the vehicle may transmit to the server data indicating a temporary deviation from the predetermined trajectory, when the actual trajectory is modified. The data may indicate a cause of the deviation, or the server may analyze the data to determine a cause of the

deviation. Knowing the cause of the deviation may be useful. For example, when the deviation is due to the driver noticing an accident that has recently occurred and, in response, steering the wheel to avoid collision, the server may plan a more moderate adjustment to the model or a specific trajectory associated with the road segment based on the cause of deviation. As another example, when the cause of deviation is a pedestrian crossing the road, the server may determine that there is no need to change the trajectory in the future.

[0341] By way of further example, the environmental image may include a lane marking indicating that vehicle 1205 is driving slightly outside of a lane, perhaps under the control of a human driver. The processor may detect the lane marking from the captured images and may modify the actual trajectory of vehicle 1205 to account for the departure from the lane. For example, a translation may be applied to the reconstructed trajectory so that it falls within the center of an observed lane.

[0342] Distributing Crowdsourced Sparse Maps

[0343] The disclosed systems and methods may enable autonomous vehicle navigation (e.g., steering control) with low footprint maps, which may be collected by the autonomous vehicles themselves without the aid of expensive surveying equipment. To support the autonomous navigation (e.g., steering applications), the road model may include a sparse map having the geometry of the road, its lane structure, and landmarks that may be used to determine the location or position of vehicles along a trajectory included in the model. As discussed above, generation of the sparse map may be performed by a remote server that communicates with vehicles travelling on the road and that receives data from the vehicles. The data may include sensed data, trajectories reconstructed based on the sensed data, and/or recommended trajectories that may represent modified reconstructed trajectories. As discussed below, the server may transmit the model back to the vehicles or other vehicles that later travel on the road to aid in autonomous navigation.

[0344] FIG. 20 illustrates a block diagram of server 1230. Server 1230 may include a communication unit 2005, which may include both hardware components (e.g., communication control circuits, switches, and antenna), and software components (e.g., communication protocols, computer codes). For example, communication unit 2005 may include at least one network interface. Server 1230 may communicate with vehicles 1205, 1210, 1215, 1220, and 1225 through communication unit 2005. For example, server 1230 may receive, through communication unit 2005, navigation information transmitted from vehicles 1205, 1210, 1215, 1220, and 1225. Server 1230 may distribute, through communication unit 2005, the autonomous vehicle road navigation model to one or more autonomous vehicles.

[0345] Server 1230 may include at least one non-transitory storage medium 2010, such as a hard drive, a compact disc, a tape, etc. Storage device 2010 may be configured to store data, such as navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225 and/or the autonomous vehicle road navigation model that server 1230 generates based on the navigation information. Storage device 2010 may be configured to store any other information, such as a sparse map (e.g., sparse map 800 discussed above with respect to FIG. 8).

[0346] In addition to or in place of storage device 2010, server 1230 may include a memory 2015. Memory 2015 may be similar to or different from memory 1140 or 150. Memory 2015 may be a non-transitory memory, such as a flash memory, a random access memory, etc. Memory 2015 may be configured to store data, such as computer codes or instructions executable by a processor (e.g., processor 2020), map data (e.g., data of sparse map 3800), the autonomous vehicle road navigation model, and/or navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225.

[0347] Server 1230 may include at least one processing device 2020 configured to execute computer codes or instructions stored in memory 2015 to perform various functions. For example, processing device 2020 may analyze the navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225, and generate the autonomous vehicle road navigation model based on the analysis. Processing device 2020 may control communication unit 1140 to distribute the autonomous vehicle road navigation model to one or more autonomous vehicles (e.g., one or more of vehicles 1205, 1210, 1215, 1220, and 1225 or any vehicle that travels on road segment 1200 at a later time). Processing device 2020 may be similar to or different from processor 1180, 1190, or processing unit 1110.

[0348] FIG. 21 illustrates a block diagram of memory 2015, which may store computer code or instructions for performing one or more operations for generating a road navigation model for use in autonomous vehicle navigation. As shown in FIG. 21, memory 2015 may store one or more modules for performing the operations for processing vehicle navigation information. For example, memory 2015 may include a model generating module 2105 and a model distributing module 2110. Processor 2020 may execute the instructions stored in any of modules 2105 and 2110 included in memory 2015.

[0349] Model generating module 2105 may store instructions which, when executed by processor 2020, may generate at least a portion of an autonomous vehicle road navigation model for a common road segment (e.g., road segment 1200) based on navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225. For example, in generating the autonomous vehicle road navigation model, processor 2020 may cluster vehicle trajectories along the common road segment 1200 into different clusters. Processor 2020 may determine a target trajectory along the common road segment 1200 based on the clustered vehicle trajectories for each of the different clusters. Such an operation may include finding a mean or average trajectory of the clustered vehicle trajectories (e.g., by averaging data representing the clustered vehicle trajectories) in each cluster. In some embodiments, the target trajectory may be associated with a single lane of the common road segment 1200.

[0350] The autonomous vehicle road navigation model may include a plurality of target trajectories each associated with a separate lane of the common road segment 1200. In some embodiments, the target trajectory may be associated with the common road segment 1200 instead of a single lane of the road segment 1200. The target trajectory may be represented by a three-dimensional spline. In some embodiments, the spline may be defined by less than 110 kilobytes per kilometer, less than 20 kilobytes per kilometer, less than 100 kilobytes per kilometer, less than 1 megabyte per kilometer, or any other suitable storage size per kilometer. Model distributing module 2110 may then distribute the generated model to one or more vehicle, e.g., as discussed below with respect to FIG. 24.

[0351] The road model and/or sparse map may store trajectories associated with a road segment. These trajectories may be referred to as target trajectories, which are provided to autonomous vehicles for autonomous navigation. The target trajectories may be received from multiple vehicles, or may be generated based on actual trajectories or recommended trajectories (actual trajectories with some modifications) received from multiple vehicles. The target trajectories included in the road model or sparse map may be continuously updated (e.g., averaged) with new trajectories received from other vehicles.

[0352] Vehicles travelling on a road segment may collect data by various sensors. The data may include landmarks, road signature profile, vehicle motion (e.g., accelerometer data, speed data), vehicle position (e.g., GPS data), and may either reconstruct the actual trajectories themselves, or transmit the data to a server, which will reconstruct the actual trajectories for the vehicles. In some embodiments, the vehicles may transmit data relating to a trajectory (e.g., a curve in an arbitrary reference frame), landmarks data, and lane assignment along traveling path to server 1230. Various vehicles travelling along the same road segment at multiple drives may have different trajectories. Server 1230 may identify routes or trajectories associated with each lane from the trajectories received from vehicles through a clustering process.

[0353] FIG. 22 illustrates a process of clustering vehicle trajectories associated with vehicles 1205, 1210, 1215, 1220, and 1225 for determining a target trajectory for the common road segment (e.g., road segment 1200). The target trajectory or a plurality of target trajectories determined from the clustering process may be included in the autonomous vehicle road navigation model or sparse map 800. In some embodiments, vehicles 1205, 1210, 1215, 1220, and 1225 traveling along road segment 1200 may transmit a plurality of trajectories 2200 to server 1230. In some embodiments, server 1230 may generate trajectories based on landmark, road geometry, and vehicle motion information received from vehicles 1205, 1210, 1215, 1220, and 1225. To generate the autonomous vehicle road navigation model, server 1230 may cluster vehicle trajectories 1600 into a plurality of clusters 2205, 2210, 2215, 2220, 2225, and 2230, as shown in FIG. 22.

[0354] Clustering may be performed using various criteria. In some embodiments, all drives in a cluster may be similar with respect to the absolute heading along the road segment 1200. The absolute heading may be obtained from GPS signals received by vehicles 1205, 1210, 1215, 1220, and 1225. In some embodiments, the absolute heading may be obtained using dead reckoning. Dead reckoning, as one of skill in the art would understand, may be used to determine the current position and hence heading of vehicles 1205, 1210, 1215, 1220, and 1225 by using previously determined position, estimated speed, etc. Trajectories clustered by absolute heading may be useful for identifying routes along the roadways.

[0355] In some embodiments, all the drives in a cluster may be similar with respect to the lane assignment (e.g., in the same lane before and after a junction) along the drive on road segment 1200. Trajectories clustered by lane assignment may be useful for identifying lanes along the roadways. In some embodiments, both criteria (e.g., absolute heading and lane assignment) may be used for clustering.

[0356] In each cluster 2205, 2210, 2215, 2220, 2225, and 2230, trajectories may be averaged to obtain a target trajectory associated with the specific cluster. For example, the trajectories from multiple drives associated with the same lane cluster may be averaged. The averaged trajectory may be a target trajectory associated with a specific lane. To average a cluster of trajectories, server 1230 may select a reference frame of an arbitrary trajectory CO. For all other trajectories (C1, ..., Cn), server 1230 may find a rigid transformation that maps Ci to CO, where $i=1, 2, \dots, n$, where n is a positive integer number, corresponding to the total number of trajectories included in the cluster. Server 1230 may compute a mean curve or trajectory in the CO reference frame.

[0357] In some embodiments, the landmarks may define an arc length matching between different drives, which may be used for alignment of trajectories with lanes. In some embodiments, lane marks before and after a junction may be used for alignment of trajectories with lanes.

[0358] To assemble lanes from the trajectories, server 1230 may select a reference frame of an arbitrary lane. Server 1230 may map partially overlapping lanes to the selected reference frame. Server 1230 may continue mapping until all lanes are in the same reference frame. Lanes that are next to each other may be aligned as if they were the same lane, and later they may be shifted laterally.

[0359] Landmarks recognized along the road segment may be mapped to the common reference frame, first at the lane level, then at the junction level. For example, the same landmarks may be recognized multiple times by multiple vehicles in multiple drives. The data regarding the same landmarks received in different drives may be slightly different. Such data may be averaged and mapped to the same reference frame, such as the CO reference frame. Additionally or alternatively, the variance of the data of the same landmark received in multiple drives may be calculated.

[0360] In some embodiments, each lane of road segment 120 may be associated with a target trajectory and certain landmarks. The target trajectory or a plurality of such target trajectories may be included in the autonomous vehicle road navigation model, which may be used later by other autonomous vehicles travelling along the same road segment 1200. Landmarks identified by vehicles 1205, 1210, 1215, 1220, and 1225 while the vehicles travel along road segment 1200 may be recorded in association with the target trajectory. The data of the target trajectories and landmarks may be continuously or periodically updated with new data received from other vehicles in subsequent drives.

[0361] For localization of an autonomous vehicle, the disclosed systems and methods may use an Extended Kalman Filter. The location of the vehicle may be determined based on three-dimensional position data and/or three-dimensional orientation data, prediction of future location ahead of vehicle's current location by integration of ego motion. The localization of vehicle may be corrected or adjusted by image observations of landmarks. For example, when vehicle detects a landmark within an image captured by the camera, the landmark may be compared to a known landmark stored within the road model or sparse map 800. The known landmark may have a known location (e.g., GPS data) along a target trajectory stored in the road model and/or sparse map 800. Based on the current speed and images of the landmark, the distance from the vehicle to the landmark may be estimated. The location of the vehicle along a target trajectory may be adjusted based on the distance to the landmark and the

landmark's known location (stored in the road model or sparse map 800). The landmark's position/location data (e.g., mean values from multiple drives) stored in the road model and/or sparse map 800 may be presumed to be accurate.

[0362] In some embodiments, the disclosed system may form a closed loop subsystem, in which estimation of the vehicle's six degrees of freedom location (e.g., three dimensional position data plus three dimensional orientation data) may be used for navigating (e.g., steering the wheel of) the autonomous vehicle to reach a desired point (e.g., 1.33 seconds ahead in the stored). In turn, data measured from the steering and actual navigation may be used to estimate the six degrees of freedom location.

[0363] In some embodiments, poles along a road, such as lampposts and power or cable line poles, may be used as landmarks for localizing the vehicles. Other landmarks such as traffic signs, traffic lights, arrows on the road, stop lines, as well as static features or signatures of an object along the road segment may also be used as landmarks for localizing the vehicle. When poles are used for localization, the x observation of the poles (i.e., the viewing angle from the vehicle) may be used, rather than the y observation (i.e., the distance to the pole) since the bottoms of the poles may be occluded and sometimes they are not on the road plane.

[0364] FIG. 23 illustrates a navigation system for a vehicle, which may be used for autonomous navigation using a crowdsourced sparse map. For illustration, the vehicle is referenced as vehicle 1205. The vehicle shown in FIG. 23 may be any other vehicle disclosed herein, including, for example, vehicles 1210, 1215, 1220, and 1225, as well as vehicle 200 shown in other embodiments. As shown in FIG. 12, vehicle 1205 may communicate with server 1230. Vehicle 1205 may include an image capture device 122 (e.g., camera 122). Vehicle 1205 may include a navigation system 2300 configured for providing navigation guidance for vehicle 1205 to travel on a road (e.g., road segment 1200). Vehicle 1205 may also include other sensors, such as a speed sensor 2320 and an accelerometer 2325. Speed sensor 2320 may be configured to detect the speed of vehicle 1205. Accelerometer 2325 may be configured to detect an acceleration or deceleration of vehicle 1205. Vehicle 1205 shown in FIG. 23 may be an autonomous vehicle, and the navigation system 2300 may be used for providing navigation guidance for autonomous driving. Alternatively, vehicle 1205 may also be a non-autonomous, human-controlled vehicle, and navigation system 2300 may still be used for providing navigation guidance.

[0365] Navigation system 2300 may include a communication unit 2305 configured to communicate with server 1230 through communication path 1235. Navigation system 2300 may also include a GPS unit 2310 configured to receive and process GPS signals. Navigation system 2300 may further include at least one processor 2315 configured to process data, such as GPS signals, map data from sparse map 800 (which may be stored on a storage device provided onboard vehicle 1205 and/or received from server 1230), road geometry sensed by a road profile sensor 2330, images captured by camera 122, and/or autonomous vehicle road navigation model received from server 1230. The road profile sensor 2330 may include different types of devices for measuring different types of road profile, such as road surface roughness, road width, road elevation, road curvature, etc. For example, the road profile sensor 2330 may include a device that measures the motion of a suspension of vehicle 1205 to

derive the road roughness profile. In some embodiments, the road profile sensor 2330 may include radar sensors to measure the distance from vehicle 1205 to road sides (e.g., barrier on the road sides), thereby measuring the width of the road. In some embodiments, the road profile sensor 2330 may include a device configured for measuring the up and down elevation of the road. In some embodiment, the road profile sensor 2330 may include a device configured to measure the road curvature. For example, a camera (e.g., camera 122 or another camera) may be used to capture images of the road showing road curvatures. Vehicle 1205 may use such images to detect road curvatures.

[0366] The at least one processor 2315 may be programmed to receive, from camera 122, at least one environmental image associated with vehicle 1205. The at least one processor 2315 may analyze the at least one environmental image to determine navigation information related to the vehicle 1205. The navigation information may include a trajectory related to the travel of vehicle 1205 along road segment 1200. The at least one processor 2315 may determine the trajectory based on motions of camera 122 (and hence the vehicle), such as three dimensional translation and three dimensional rotational motions. In some embodiments, the at least one processor 2315 may determine the translation and rotational motions of camera 122 based on analysis of a plurality of images acquired by camera 122. In some embodiments, the navigation information may include lane assignment information (e.g., in which lane vehicle 1205 is travelling along road segment 1200). The navigation information transmitted from vehicle 1205 to server 1230 may be used by server 1230 to generate and/or update an autonomous vehicle road navigation model, which may be transmitted back from server 1230 to vehicle 1205 for providing autonomous navigation guidance for vehicle 1205.

[0367] The at least one processor 2315 may also be programmed to transmit the navigation information from vehicle 1205 to server 1230. In some embodiments, the navigation information may be transmitted to server 1230 along with road information. The road location information may include at least one of the GPS signal received by the GPS unit 2310, landmark information, road geometry, lane information, etc. The at least one processor 2315 may receive, from server 1230, the autonomous vehicle road navigation model or a portion of the model. The autonomous vehicle road navigation model received from server 1230 may include at least one update based on the navigation information transmitted from vehicle 1205 to server 1230. The portion of the model transmitted from server 1230 to vehicle 1205 may include an updated portion of the model. The at least one processor 2315 may cause at least one navigational maneuver (e.g., steering, such as making a turn, braking, accelerating, passing another vehicle, etc.) by vehicle 1205 based on the received autonomous vehicle road navigation model or the updated portion of the model.

[0368] The at least one processor 2315 may be configured to communicate with various sensors and components included in vehicle 1205, including communication unit 705, GPS unit 2315, camera 122, speed sensor 2320, accelerometer 2325, and road profile sensor 2330. The at least one processor 2315 may collect information or data from various sensors and components, and transmit the information or data to server 1230 through communication unit 2305. Alternatively or additionally, various sensors or

components of vehicle 1205 may also communicate with server 1230 and transmit data or information collected by the sensors or components to server 1230.

[0369] In some embodiments, vehicles 1205, 1210, 1215, 1220, and 1225 may communicate with each other, and may share navigation information with each other, such that at least one of the vehicles 1205, 1210, 1215, 1220, and 1225 may generate the autonomous vehicle road navigation model using crowdsourcing, e.g., based on information shared by other vehicles. In some embodiments, vehicles 1205, 1210, 1215, 1220, and 1225 may share navigation information with each other and each vehicle may update its own the autonomous vehicle road navigation model provided in the vehicle. In some embodiments, at least one of the vehicles 1205, 1210, 1215, 1220, and 1225 (e.g., vehicle 1205) may function as a hub vehicle. The at least one processor 2315 of the hub vehicle (e.g., vehicle 1205) may perform some or all of the functions performed by server 1230. For example, the at least one processor 2315 of the hub vehicle may communicate with other vehicles and receive navigation information from other vehicles. The at least one processor 2315 of the hub vehicle may generate the autonomous vehicle road navigation model or an update to the model based on the shared information received from other vehicles. The at least one processor 2315 of the hub vehicle may transmit the autonomous vehicle road navigation model or the update to the model to other vehicles for providing autonomous navigation guidance.

[0370] FIG. 24 is a flowchart showing an exemplary process 2400 for generating a road navigation model for use in autonomous vehicle navigation. Process 2400 may be performed by server 1230 or processor 2315 included in a hub vehicle. In some embodiments, process 2400 may be used for aggregating vehicle navigation information to provide an autonomous vehicle road navigation model or to update the model.

[0371] Process 2400 may include receiving, by a server, navigation information from a plurality of vehicles (step 2405). For example, server 1230 may receive the navigation information from vehicles 1205, 1210, 1215, 1220, and 1225. The navigation information from the plurality of vehicles may be associated with a common road segment (e.g., road segment 1200) along which the plurality of vehicles, e.g., 1205, 1210, 1215, 1220, and 1225, travel.

[0372] Process 2400 may further include storing, by the server, the navigation information associated with the common road segment (step 2410). For example, server 1230 may store the navigation information in storage device 2010 and/or memory 2015.

[0373] Process 2400 may further include generating, by the server, at least a portion of an autonomous vehicle road navigation model for the common road segment based on the navigation information from the plurality of vehicles (step 2415). The autonomous vehicle road navigation model for the common road segment may include at least one line representation of a road surface feature extending along the common road segment, and each line representation may represent a path along the common road segment substantially corresponding with the road surface feature. For example, the road surface feature may include a road edge or a lane marking. Moreover, the road surface feature may be identified through image analysis of a plurality of images acquired as the plurality of vehicles traverse the

common road segment. For example, server 1230 may generate at least a portion of the autonomous vehicle road navigation model for common road segment 1200 based on the navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225 that travel on the common road segment 1200.

[0374] In some embodiments, the autonomous vehicle road navigation model may be configured to be superimposed over a map, an image, or a satellite image. For example, the model may be superimposed over a map or image provided by a conventional navigation service such as Google Maps, Waze, or the like.

[0375] In some embodiments, generating at least a portion of the autonomous vehicle road navigation model may include identifying, based on image analysis of the plurality of images, a plurality of landmarks associated with the common road segment. In certain aspects, this analysis may include accepting potential landmarks when a ratio of images in which the landmark does appear to images in which the landmark does not appear exceeds a threshold and/or rejecting potential landmarks when a ratio of images in which the landmark does not appear to images in which the landmark does appear exceeds a threshold. For example, if a potential landmark appears in data from vehicle 1210 but not in data from vehicles 1205, 1215, 1220, and 1225, the system may determine that a ratio of 1:5 is below the threshold for accepting the potential landmark. By way of further example, if a potential landmark appears in data from vehicles 1205, 1215, 1220, and 1225 but not in data from vehicle 1210, the system may determine that a ratio of 4:5 is above the threshold for accepting the potential landmark.

[0376] Process 2400 may further include distributing, by the server, the autonomous vehicle road navigation model to one or more autonomous vehicles for use in autonomously navigating the one or more autonomous vehicles along the common road segment (step 2420). For example, server 1230 may distribute the autonomous vehicle road navigation model or a portion (e.g., an update) of the model to vehicles 1205, 1210, 1215, 1220, and 1225, or any other vehicles later travel on road segment 1200 for use in autonomously navigating the vehicles along road segment 1200.

[0377] Process 2400 may include additional operations or steps. For example, generating the autonomous vehicle road navigation model may include clustering vehicle trajectories received from vehicles 1205, 1210, 1215, 1220, and 1225 along road segment 1200 into a plurality of clusters and/or aligning data received from vehicles 1205, 1210, 1215, 1220, and 1225, as discussed in further detail below, with respect to FIG. 29. Process 2400 may include determining a target trajectory along common road segment 1200 by averaging the clustered vehicle trajectories in each cluster. Process 2400 may also include associating the target trajectory with a single lane of common road segment 1200. Process 2400 may include determining a three-dimensional spline to represent the target trajectory in the autonomous vehicle road navigation model.

[0378] Using Crowdsourced Sparse Maps for Navigation

[0379] As discussed above, server 1230 may distribute a generated road navigation model to one or more vehicles. The road navigation model may be included in a sparse map, as described in detail above. Consistent with embodiments of the present disclosure, one or more vehicles may be configured to use the distributed sparse map for autonomous navigation.

[0380] FIG. 25 is an exemplary functional block diagram of memory 1140 and/or 1150, which may be stored/programmed with instructions for performing one or more operations consistent with the disclosed embodiments. Although the following refers to memory 1140, one of skill in the art will recognize that instructions may be stored in memory 1140 and/or 1150.

[0381] As shown in FIG. 25, memory 1140 may store a sparse map module 2502, an image analysis module 2504, a road surface feature module 2506, and a navigational response module 2508. The disclosed embodiments are not limited to any particular configuration of memory 1140. Further, applications processor 1180 and/or image processor 1190 may execute the instructions stored in any of modules 2502, 2504, 2506, and 2508 included in memory 1140. One of skill in the art will understand that references in the following discussions to processing unit 110 may refer to applications processor 1180 and image processor 1190 individually or collectively. Accordingly, steps of any of the following processes may be performed by one or more processing devices.

[0382] In one embodiment, sparse map module 2502 may store instructions which, when executed by processing unit 110, receive (and, in some embodiments, store) a sparse map distributed by server 1230. Sparse map module 2502 may receive an entire sparse map in one communication or may receive a sub-portion of a sparse map, the sub-portion corresponding to an area in which the vehicle is operating.

[0383] In one embodiment, image analysis module 2504 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs analysis of one or more images acquired by one of image capture devices 122, 124, and 126. As described in further detail below, image analysis module 2504 may analyze the one or more images to determine a current position of the vehicle.

[0384] In one embodiment, road surface feature module 2506 may store instructions which, when executed by processing unit 110, identifies a road surface feature in the sparse map received by sparse map module 2502 and/or in the one or more images acquired by one of image capture devices 122, 124, and 126.

[0385] In one embodiment, navigational response module 2508 may store software executable by processing unit 110 to determine a desired navigational response based on data derived from execution of sparse map module 2502, image analysis module 2504, and/or road surface feature module 2506.

[0386] Furthermore, any of the modules (e.g., modules 2502, 2504, and 2506) disclosed herein may implement techniques associated with a trained system (such as a neural network or a deep neural network) or an untrained system.

[0387] FIG. 26 is a flowchart showing an exemplary process 2600 for autonomously navigating a vehicle along a road segment. Process 2600 may be performed by processor 2315 included in navigation system 2300.

[0388] Process 2600 may include receiving a sparse map model (step 2605). For example, processor 2315 may receive the sparse map from server 1230. In some embodiments, the sparse map model may include at least one line representation of a road surface feature extending along the road.

segment, and each line representation may represent a path along the road segment substantially corresponding with the road surface feature. For example, the road feature may include a road edge or a lane marking.

[0389] Process 2600 may further include receiving, from a camera, at least one image representative of an environment of the vehicle (step 2610). For example, processor 2315 may receive, from camera 1122, the at least one image. (Camera 1122 may capture one or more images of the environment surrounding vehicle 11205 as vehicle 11205 travels along road segment 11200.

[0390] Process 2600 may also include analyzing the sparse map model and the at least one image received from the camera (step 2615). For example, analysis of the sparse map model and the at least one image received from the camera may include determining a current position of the vehicle relative to a longitudinal position along the at least one line representation of a road surface feature extending along the road segment. In some embodiments, this determination may be based on identification of at least one recognized landmark in the at least one image. In some embodiments, process 2600 may further include determining an estimated offset based on an expected position of the vehicle relative to the longitudinal position and the current position of the vehicle relative to the longitudinal position.

[0391] Process 2600 may further include determining an autonomous navigational response for the vehicle based on the analysis of the sparse map model and the at least one image received from the camera (step 2620). In embodiments in which processor 2315 determines an estimated offset, the autonomous navigational response may be further based on the estimated offset. For example, if processor 2315 determines that the vehicle is offset from the at least one line representation by 1 m to the left, processor 2315 may cause the vehicle to shift towards the right (e.g., by changing an orientation of the wheels). By way of further example, if processor 2315 determines that an identified landmark is offset from an expected position, processor 2315 may cause the vehicle to shift so as to move the identified landmark towards its expected position. Accordingly, in some embodiments, process 2600 may further include adjusting a steering system of the vehicle based on the autonomous navigational response.

[0392] Aliening Crowdsourced Map Data

[0393] As discussed above, generation of a crowdsourced sparse map may use data from a plurality of drives along a common road segment. This data may be aligned in order to generate a coherent sparse map. As discussed above with respect to FIG. 14, generating a map skeleton may be insufficient to construct splines for use in navigation. Thus, embodiments of the present disclosure may allow for aligning data crowdsourced from a plurality of drives.

[0394] FIG. 27 illustrates a block diagram of memory 2015, which may store computer code or instructions for performing one or more operations for generating a road navigation model for use in autonomous vehicle navigation. As shown in FIG. 21, memory 2015 may store one or more modules for performing the operations for processing vehicle navigation information. For example, memory 2015 may include a drive data receiving module 2705 and a longitudinal alignment module 2710. Processor 2020 may execute the instructions stored in any of modules 2705 and 2710 included in memory 2015.

[0395] Drive data receiving module 2705 may store instructions which, when executed by processor 2020, may control communication device 2005 to receive drive data from one or more vehicles (e.g., 1205, 1210, 1215, 1220, and 1225).

[0396] Longitudinal alignment module 2710 may store instructions which, when executed by processor 2020, align the data received using drive data receiving module 2705 when the data is related to common road segment (e.g., road segment 1200) based on navigation information received from vehicles 1205, 1210, 1215, 1220, and 1225. For example, longitudinal alignment module 2710 may align the data along patches, which may allow for easier optimization of the error correction that alignment entails. In some embodiments, longitudinal alignment module 2710 may further score each patch alignment with a confidence score.

[0397] FIG. 28A illustrates an example of raw location data from four different drives. In the example of FIG. 28A, the raw data from the first drive is depicted as a series of stars, the raw data from the second drive is depicted as a series of filled-in squares, the raw data from third drive is depicted as a series of open squares, and the raw data from the fourth drive is depicted as a series of open circles. As one skilled in the art would recognize, the shapes are merely illustrative of the data itself, which may be stored as a series of coordinates, whether local or global.

[0398] As may be seen in FIG. 28A, the drives may occur in different lanes along the same road (represented by line 1200). Furthermore, FIG. 28A depicts that the drive data may include variances due to errors in locational measurements (e.g., GPS) and may have missing data points due to system errors. Finally, FIG. 28A also depicts that each drive may start and end at a different point within a segment along the road.

[0399] FIG. 28B illustrates another example of raw location data from five different drives. In the example of FIG. 28B, the raw data from the first drive is depicted as a series of filled-in squares, the raw data from the second drive is depicted as a series of open squares, the raw data from third drive is depicted as a series of open circles, the raw data from the fourth drive is depicted as a series of stars, and the raw data from the fifth drive is depicted as a series of triangles. As one skilled in the art would recognize, the shapes are merely illustrative of the data itself, which may be stored as a series of coordinates, whether local or global.

[0400] FIG. 28B illustrates similar properties of the drive data as FIG. 28A. FIG. 28B further depicts that an intersection may be detected by tracking the movement of the fifth drive away from the others. For example, the example data in FIG. 28B may suggest that an exit ramp is present on the right side of the road (represented by line 1200). FIG. 28B also depicts that added lanes may be detected if data begins on a new portion of the road. For example, the fourth drive in the example data in FIG. 28B may suggest that a fourth lane is added to the road shortly after the detected exit ramp.

[0401] FIG. 28C illustrates an example of raw location data with target trajectories therefrom. For example, the first drive data (represented by triangles) and the second drive data (represented by open squares) have an associated target trajectory 2810. Similarly, the third drive data (represented by open

circles) has an associated target trajectory 2820, and the fourth drive data (represented by filled-in squares) has an associated target trajectory 2830.

[0402] In some embodiments, drive data may be reconstructed such that one target trajectory is associated with each lane of travel, as depicted in FIG. 28C. Such target trajectories may be generated from one or more simple, smooth line models if proper alignment of the patches of the drive data is performed. Process 2900, discussed below, is one example of proper alignment of patches.

[0403] FIG. 29 is a flowchart showing an exemplary process 2900 for determining a line representation of a road surface feature extending along a road segment. The line representation of the road surface feature may be configured for use in autonomous vehicle navigation, e.g., using process 2600 of FIG. 26 above. Process 2900 may be performed by server 1230 or processor 2315 included in a hub vehicle.

[0404] Process 2900 may include receiving, by a server, a first set of drive data including position information associated with the road surface feature (step 2905). The position information may be determined based on analysis of images of the road segment, and the road surface feature may include a road edge or a lane marking.

[0405] Process 2900 may further include receiving, by a server, a second set of drive data including position information associated with the road surface feature (step 2910). As with step 2905, the position information may be determined based on analysis of images of the road segment, and the road surface feature may include a road edge or a lane marking. Steps 2905 and 2910 may be performed at concurrent times, or there may be a lapse of time between step 2905 and step 2910, depending on when the first set of drive data and the second set of drive data are collected.

[0406] Process 2900 may also include segmenting the first set of drive data into first drive patches and segmenting the second set of drive data into second drive patches (step 2915). A patch may be defined by a data size or by a drive length. The size or length that defines a patch may be predefined to one or more values or may be updated with a neural network or other machine technology. For example, the length of a patch may be preset to always be 1 km, or it may be preset to be 1 km when traveling along a road at greater than 30 km/hr and to be 0.8 km when traveling along a road at less than 30 km/hr. Alternatively, or concurrently, machine learning analysis may optimize the size or length of patches based on a number of dependent variables such as driving conditions, driving speed, and the like. Still further, in some embodiments, a path may be defined by a data size and a drive length.

[0407] In some embodiments, the first set of data and the second set of data may include position information and may be associated with a plurality of landmarks. In such embodiments, process 2900 may further include determining whether to accept or reject landmarks within the sets of data based on one or more thresholds, as described above with respect to FIG. 19.

[0408] Process 2900 may include longitudinally aligning the first set of drive data with the second set of drive data within corresponding patches (step 2920). For example, longitudinal alignment may include selecting either the first set of data or the second set of data as the reference data set and then shifting and/or elastically stretching the other set of data to align patches within the sets. In some

embodiments, aligning the sets of data may further include aligning GPS data included in both sets and associated with the patches. For example, the connections between patches with the sets may be adjusted to align more closely with GPS data. In such embodiments, however, the adjustment must be limited to prevent the limitations of GPS data from corrupting the alignment. For example, GPS data is not three-dimensional and, therefore, may create unnatural twists and slopes when projected onto a three-dimensional representation of the road.

[0409] Process 2900 may further include determining the line representation of the road surface feature based on the longitudinally aligned first and second drive data in the first and second draft patches (step 2925). For example, the line representation may be constructed using a smooth line model on the aligned data. In some embodiments, determining the line representation may include an alignment of the line representation with global coordinates based on GPS data acquired as part of at least one of the first set of drive data or the second set of drive data. For example, the first set of data and/or the second set of data may be represented in local coordinates; however, the use of a smooth line model requires both sets to have the same coordinate axes. Accordingly, in certain aspects, the first set of data and/or second set of data may be adjusted to have the same coordinate axes as each other.

[0410] In some embodiments, determining the line representation may include determining and applying a set of average transformations. For example, each of the average transformations may be based on transformations determined that link data from the first set of drive data across sequential patches and the link data from the second set of drive data across sequential patches.

[0411] Process 2900 may include additional operations or steps. For example, process 2900 may further include overlaying the line representation of the road surface feature on at least one geographical image. For example, the geographical image may be a satellite image. By way of further example, process 2900 may further include filtering out landmark information and/or drive data that appears erroneous based on the determined line representation and longitudinal alignment. Such filtering may be similar in concept to the rejection of possible landmarks based on one or more thresholds, as described above with respect to FIG. 19.

[0412] Crowdsourcing Road Surface Information

[0413] In addition to crowdsourcing landmarks and line representations in order to generate a sparse map, the disclosed systems and methods may crowdsource road surface information as well.

Accordingly, road conditions may be stored along with and/or within a sparse map used for navigation of an autonomous vehicle.

[0414] FIG. 30 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.

[0415] As shown in FIG. 30, memory 140 may store an image receiving module 3002, a road surface feature module 3004, a location determination module 3006, and a navigational response module 3008. The disclosed embodiments are not limited to any particular configuration of memory 140.

Further, applications processor 880 and/or image processor 190 may execute the instructions stored in any of modules 3002, 3004, 3006, and 3008 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

5 processes may be performed by one or more processing devices.

[0416] In one embodiment, image receiving module 3002 may store instructions which, when executed by processing unit 110, receive (and, in some embodiments, store) images acquired by one of image capture devices 122, 124, and 126.

10 [0417] In one embodiment, road surface feature module 3004 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs analysis of one or more images acquired by one of image capture devices 122, 124, and 126 and identifies a road surface feature. For example, the road surface feature may include a road edge or a lane marking.

15 [0418] In one embodiment, location determination module 3006 may store instructions (such as GPS software or visual odometry software) which, when executed by processing unit 110, receives location information relating to the vehicle. For example, location determination module 3006 may receive GPS data and/or ego-motion data including the position of the vehicle. In some embodiments, location determination module 3006 may calculate one or more locations using received information. For example, location determination module 3006 may receive one or more images acquired by one of image capture devices 122, 124, and 126 and determine a location of the vehicle using analysis of the images.

20 [0419] In one embodiment, navigational response module 3008 may store software executable by processing unit 110 to determine a desired navigational response based on data derived from execution of image receiving module 3002, road surface feature module 3004, and/or location determination module 3006.

25 [0420] Furthermore, any of the modules (e.g., modules 3002, 3004, and 3006) disclosed herein may implement techniques associated with a trained system (such as a neural network or a deep neural network) or an untrained system.

[0421] FIG. 3 is a flowchart showing an exemplary process 3100 for collecting road surface information for a road segment. Process 3100 may be performed by processor 2315 included in navigation system 2300.

30 [0422] Process 3100 may include receiving, from a camera, at least one image representative of a portion of the road segment (step 3105). For example, processor 2315 may receive, from camera 122, the at least one image. Camera 122 may capture one or more images of the environment surrounding vehicle 1205 as vehicle 1205 travels along road segment 1200.

35 [0423] Process 3100 may further include identifying in the at least one image at least one road surface feature along the portion of the road segment (step 3010). For example, the at least one road surface feature may include a road edge or may include a lane marking.

[0424] Process 3100 may also include determining a plurality of locations associated with the road surface feature according to a local coordinate system of the vehicle (step 3115). For example, processor 2315 may use ego-motion data and/or GPS data to determine the plurality of locations.

[0425] Process 3100 may further include transmitting the determined plurality of locations from the vehicle to a server (step 3120). For example, the determined locations may be configured to enable determination by the server of a line representation of the road surface feature extending along the road segment, as described above with respect to FIG. 229. In some embodiments, the line representation may represent a path along the road segment substantially corresponding with the road surface feature.

[0426] Process 3100 may include additional operations or steps. For example, process 3100 may further include receiving from the server the line representation. In this example, processor 2315 may receive the line representation as a portion of a sparse map received, for example, according to process 2400 of FIG. 24 and/or process 2600 of FIG. 26. By way of further example, process 2900 may further include overlaying the line representation of the road surface feature on at least one geographical image. For example, the geographical image may be a satellite image.

[0427] As explained above with respect to crowdsourcing of landmarks, in some embodiments, the server may implement selection criteria for deciding whether to accept or reject possible road surface features received from the vehicles. For example, the server may accept road surface features when a ratio of location sets in which the road surface feature does appear to location sets in which the road surface feature does not appear exceeds a threshold and/or rejecting potential road surface features when a ratio of location sets in which the road surface feature does not appear to location sets in which the road surface feature does appear exceeds a threshold.

[0428] Vehicle Localization

[0429] In some embodiments, the disclosed systems and methods may use a sparse map for autonomous vehicle navigation. In particular, the sparse map may be for autonomous vehicle navigation along a road segment. For example, the sparse map may provide sufficient information for navigating an autonomous vehicle without storing and/or updating a large quantity of data. As discussed below in further detail, an autonomous vehicle may use the sparse map to navigate one or more roads based on one or more stored trajectories.

[0430] Autonomous Vehicle Lane Localization Using Lane Markings

[0431] As explained above, a self-driving vehicle may navigate based on dead reckoning between landmarks. However, errors may accumulate during navigation by dead reckoning, and thus over time the position determinations relative to the target trajectory may become increasingly less accurate. As explained below, lane markings may be used for localization of the vehicle during landmark spacings, which may minimize accumulation of errors during navigation by dead reckoning.

[0432] For example, a spline may be represented as shown in Equation 1 below:

$$B(u) = \sum_k b_k(u) P^{(k)}$$

Equation 11

[0433] In the example of Equation 1, $B(u)$ is the curve representing the spline, $b_k(u)$ are the basis functions, and p^k represents a control point. The control point may be transformed to local coordinates according to, for example, Equation 2 below:

$$P_l^{(k)} = R^T (P^{(k)} - T)$$

Equation 2

[0434] In the example of Equation 2, $P_l^{(k)}$ represents the control point $P^{(k)}$ transformed to local coordinates, R is the rotation matrix that may, for example, be inferred from the vehicle's heading, R^T represents the transpose of the rotation matrix, and T represents the location of the vehicle.

[0435] In some embodiments, the local curve representing the path of the vehicle may be determined using Equation 3 below:

$$H_l(u) = y B_l(u) - f B_l'(u) = 0$$

Equation 3

[0436] In the example of Equation 3, f is the focal length of the camera and B_z , B_y , and B_x represent the components of curve B in local coordinates. The derivation of H may be represented by Equation 4 below:

$$H'(u) = y B_l'(u) - f B_l''(u)$$

Equation 4

[0437] Based on Equation 1, B_l' in Equation 4 may be further represented by Equation 5 below:

$$B_l'(u) = \sum_k b'_k(u) P_l^{(k)}$$

Equation 5

[0438] Equation 5 may, for example, be solved by a Newton-Raphson-based solver. In certain aspects, the solver may be run for five or fewer steps. To solve for x , in some embodiments, one may use Equation 6 below:

$$x = f \frac{B_{lx}(u_i)}{B_{lz}(u_i)}$$

Equation 6

[0439] In some embodiments, the derivatives of a trajectory may be used. For example, the derivatives may be given by Equation 7 below:

$$\frac{dx}{dX_j} = f \frac{1}{B_{lz}(u_i)} \frac{dB_{lx}(u_i)}{dX_j} + f \frac{B_{lx}(u_i)}{B_{lz}(u_i)^2} \frac{dB_{lz}(u_i)}{dX_j}$$

Equation 7

[0440] In the example of Equation 7, X_j may represent state components, for example, for the location of the vehicle. In certain aspects, j may represent an integer between one and six.

[0441] To solve Equation 7, in some embodiments, one may use Equation 8 below:

$$\frac{dB_{lx}(u_i)}{dX_j} = \sum_k \frac{d}{dX_j} (b_k(u_i) P_{lx}^{(k)}) = \sum_k \left(b'_k(u_i) \frac{du_i}{dX_j} + b_k(u_i) \frac{dP_{lx}^{(k)}}{dX_j} \right)$$

Equation 8

[0442] To solve Equation 8, in some embodiments, one may use implicit differentiation to obtain Equations 9 and 10 below:

$$\frac{\partial B_{lz}(u_i)}{\partial X_j} = \sum_k b_k(u_i) \frac{dP_{lz}^{(k)}}{dX_j}$$

$$\frac{\partial B_{ly}(u_i)}{\partial X_j} = \sum_k b_k(u_i) \frac{dP_{ly}^{(k)}}{dX_j}$$

Equations 9 and 10

[0443] Using Equations 9 and 10, the derivatives of the trajectory may be obtained. In some embodiments, then, the Extended Kalman Filter may be used to localize lane measurements. By localizing lane measurements, as explained above, lane markings may be used to minimize accumulation of errors during navigation by dead reckoning. The use of lane markings is described in further detail below with respect to FIGS. 32-35.

[0444] FIG. 32 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.

[0445] As shown in FIG. 32, memory 140 may store a position determination module 3202, an image analysis module 3204, a distance determination module 3206, and an offset determination module 3208. 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 3202, 3204, 3206, and 3208 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.

[0446] In one embodiment, position determination module 3202 may store instructions (such as GPS software or visual odometry software) which, when executed by processing unit 110, receives location information relating to the vehicle. For example, position determination module 3202 may receive GPS data and/or ego-motion data including the position of the vehicle. In some embodiments, position determination module 3202 may calculate one or more locations using received information. For example, position determination module 3202 may receive one or more images acquired by one of image capture devices 122, 124, and 126 and determine a location of the vehicle using analysis of the images.

[0447] Position determination module 3202 may also use other navigational sensors to determine the position of the vehicle. For example, a speed sensor or accelerometer may send information to position determination module 3202 for use in calculating the position of the vehicle.

[0448] In one embodiment, image analysis module 3204 may store instructions (such as computer vision software) which, when executed by processing unit 110 performs analysis of one or more images acquired by one of image capture devices 122, 124, and 126. As described in further detail

below, image analysis module 3204 may analyze the one or more images to identify at least one lane marking.

[0449] In one embodiment, distance determination module 3206 may store instructions which, when executed by processing unit 110, performs analysis of one or more images acquired by one of image capture devices 122, 124, and 126 and determines a distance from the vehicle to a lane marking identified by image analysis module 3204.

[0450] In one embodiment, offset determination module 3208 may store software executable by processing unit 110 to determine an estimated offset of the vehicle from a road model trajectory. For example, offset determination module 3208 may use a distance determined by distance determination module 3206 to calculate the estimated offset. A desired navigational response based on data derived from execution of position determination module 3202, image analysis module 3204, distance determination module 3206, and/or offset determination module 3208 may then be determined.

[0451] Furthermore, any of the modules (e.g., modules 3202, 3204, and 3206) disclosed herein may implement techniques associated with a trained system (such as a neural network or a deep neural network) or an untrained system.

[0452] FIG. 33A illustrates an example of a vehicle navigating by dead reckoning without using lane markings. In the example of FIG. 33A, the vehicle is navigating along a trajectory 3310 but does not use lane markings (e.g., markings 3320A or 3320B) for navigating.

[0453] FIG. 33B shows the example of FIG. 33A after 350m. As depicted in FIG. 33B, the trajectory 3310 of the vehicle does not fully align with the lane markings (e.g., markings 3320A or 3320B) due to dead reckoning error accumulation.

[0454] FIG. 33C shows the example of FIGS. 33A and 33B after 1km. As depicted in FIG. 33C, the expected position 3330A of a landmark does not align with the actual position 3330B of the landmark. Although the vehicle may now use the landmark to correct the dead reckoning errors that accumulated over the course of 1km, systems and methods of the present disclosure may allow for using lane markings to minimize dead reckoning error accumulation between landmarks.

[0455] FIG. 34A illustrates an example of a vehicle navigating by dead reckoning with lane markings. In the example of FIG. 34A, the vehicle is navigating along a trajectory 3410 and also identified lane markings (e.g., markings 3420A and 3420B) to localize and use for navigating.

[0456] FIG. 34B shows the example of FIG. 34A after 350m. As depicted in FIG. 34B, the trajectory 3410 of the vehicle substantially aligns with the lane markings (e.g., markings 3420A and 3420B) because the vehicle has been correcting for dead reckoning error using the identified lane markings.

[0457] FIG. 34C shows the example of FIGS. 34A and 34B after 1km. As depicted in FIG. 34C, the expected position 3430A of a landmark substantially aligns with the actual position 3430B of the landmark. Accordingly, upon encountering the landmark, the vehicle in FIG. 34C may undertake a substantially smaller correction than the vehicle of FIG. 33C. Process 3500 of FIG. 35, discussed below, is an exemplary process by which a vehicle may use lane markings for navigation like in FIGS. 34A-34C.

[0458] FIG. 335 is a flowchart showing an exemplary process 3500 for correcting a position of a vehicle navigating a road segment. Process 3500 may be performed by processor 2315 included in a navigation system 2300.

[0459] Process 3500 may include determining, based on an output of at least one navigational sensor, a measured position of the vehicle along a predetermined road model trajectory (step 3505). For example, the predetermined road model trajectory may be associated with the road segment, and, in some embodiments, the predetermined road model trajectory may include a three-dimensional polynomial representation of a target trajectory along the road segment. The at least one navigational sensor may, for example, include a speed sensor or an accelerometer.

[0460] Process 3500 may further include receiving, from an image capture device, at least one image representative of an environment of the vehicle (step 3510). For example, processor 2315 may receive, from camera 1122, the at least one image. Camera 1122 may capture one or more images of the environment surrounding vehicle 1120 as vehicle 1120 travels along road segment 1120.

[0461] Process 3500 may also include analyzing the at least one image to identify at least one lane marking. The at least one lane marking may be associated with a lane of travel along the road segment. Process 3500 may further include determining, based on the at least one image, a distance from the vehicle to the at least one lane marking. For example, a variety of known algorithms for calculating the distance of an object in an image may be used.

[0462] Process 3500 may include determining an estimated offset of the vehicle from the predetermined road model trajectory based on the measured position of the vehicle and the determined distance. In some embodiments, determining the estimated offset may further include determining, based on the distance to the at least one lane marking, whether the vehicle is on a trajectory that will intersect the at least one lane marking. Alternatively or concurrently, in some embodiments, determining the estimated offset may further include determining, based on the distance to the at least one lane marking, whether the vehicle is within a predetermined threshold of the at least one lane marking.

[0463] Process 3500 may further include determining an autonomous steering action for the vehicle based on the estimated offset to correct the position of the vehicle. For example, in embodiments where the estimated offset includes determining whether the vehicle is on a trajectory that will intersect the at least one lane marking and/or whether the vehicle is within a predetermined threshold of the at least one lane marking, processor 2315 may determine an autonomous steering action to prevent the intersection of the vehicle trajectory with the at least one lane marking and/or to bring the distance between the vehicle and the lane marking below the predetermined threshold.

[0464] In some embodiments, determining the autonomous steering action may further include solving for at least one derivative of the predetermined road model trajectory. For example, the derivative of a trajectory may be calculated using the Extended Kalman Filter, discussed above.

[0465] Process 3500 may include additional operations or steps. For example, process 3500 may further include adjusting a steering system of the vehicle based on the autonomous steering action.

[0466] Lane Mapping and Navigation

[[0467]] The disclosed systems and embodiments may allow for lane splits to be mapped based on vehicle trajectories. The disclosed systems and embodiments may further allow for lane splits to be detected based on vehicle trajectories.

[[0468]] A lane split may include a portion of a road segment where the number of lanes increases or decreases after the split. For example, the number of available lanes for traveling in a given direction along a road segment may increase from 1 to 2, from 2 to 3, etc. A server, such as server 1230, may receive navigation information from vehicles traveling along the road segment and identify the lane split in the road segment. Server 1230 may update a vehicle road navigation model to establish target trajectories for the new lane emerging from the lane split. Accordingly, during navigation, host vehicles may enter the new lane and navigate according to the mapped trajectory.

[[0469]] The techniques for mapping and navigating lane split features in a road segment, as disclosed herein, overcome several important technical problems, including problems relating to efficiency and accuracy. Existing techniques may determine lane split features based on road boundaries or lane markings, for example, based on images captured by vehicles navigating the road segment. While in some instances, this may allow a road split to be detected, it may not reflect the most desirable path for a vehicle to travel. For example, many vehicles begin to move into a new lane before the markings on the road segment indicate the lane split feature. The techniques for mapping and navigating a lane split feature based on actual trajectories, as disclosed herein, provide target trajectories more consistent with how a vehicle would traverse the lane split feature.

[[0470]] FIG. 36 illustrates an example road segment 3600 having a lane split, consistent with the disclosed embodiments. As discussed above, a lane split may include any portion of a road in which the number of available lanes for travel increases (or decreases) along the direction of travel associated with the road. For example, as shown in FIG. 36, road segment 3600 may include lanes 3611 and 3612, in which vehicles 3621 and 3622 may be traveling, respectively. Vehicles 3620, 3621, and 3622 may be traveling along road segment 3600 in the general direction of travel indicated by arrow 3630. Road segment 3600 may include a lane split in which new lane 3610 emerges. Vehicle 3620 may be entering new lane 3610 from lane 3611, as shown in FIG. 36. As a result of the lane split, the number of available lanes of travel in road segment 3600 may increase from two lanes to three. The lane split illustrated in FIG. 36 is provided by way of example. It is to be understood that the disclosed systems and methods may apply to various other configurations of lane splits. For example, the number of available lanes of travel may increase from one to two, from two to three, from three to four, etc. In some embodiments, the number of lanes may increase by more than one (e.g., where two, three, four or more lanes are added). While the present disclosure is generally discussed in reference to an increase in the number of lanes along a given direction of travel, the disclosed systems and methods may also apply to situations in which the number of lanes decreases (for example, a lane merge, etc.).

[[0471]] Various scenarios may exist in which new lane 3610 is added as a result of a lane split. In some embodiments, the lane split may be associated with an exit from a highway or freeway. Accordingly, new lane 3610 may be associated with an exit ramp or the like. In other embodiments, the

lane split may be associated with a turn lane. For example, a vehicle 3620 may be approaching an intersection and new lane 3610 may comprise a turn lane for vehicles turning at the intersection. The lane split may also be associated with a fork or split in the road and new lane 3610 may eventually diverge from lanes 361 and/or 3612. Accordingly, new lane 3610 may eventually be associated with a separate road segment other than lanes 361 and/or 3612 after the lane split. In some embodiments, the lane split may be associated with a change in width of a stretch of road, for example, to accommodate additional traffic volume (e.g., adding a lane to a road), etc. Accordingly, new lane 3610 may continue along road segment 3600 after the split. While new lane 3610 is shown to emerge to form an outer lane of road segment 3600, it may also be an inner lane. For example, in some embodiments, new lane 3610 may emerge between lanes 3611 and 3612, etc.

[0472] Lane splits may be identified by analyzing the trajectories of vehicles traveling along road segment 3600. FIG. 37A illustrates example vehicle trajectories that may be used for identifying a lane split, consistent with the disclosed embodiments. A server, such as server 1230 may receive a plurality of vehicle trajectories 3701, 3702, 3703, 3704, 3705, and 3706 from vehicles traversing road segment 3600 and may use them to identify the lane split. Trajectories 3701, 3702, and 3703 may represent trajectories of vehicles electing to enter new lane 3610 after the lane split. Trajectories 3704, 3705, and 3706 may represent trajectories of vehicles continuing to travel within lane 3611 after the lane split. The trajectories may be recorded by the vehicles and transmitted from the vehicles to server 1230. As discussed in greater detail above, the trajectories may be represented by a spline, a series of points (e.g., location identifiers), or in various other forms. In some embodiments, the trajectories may be defined based on a reference coordinate system, as discussed above. In some embodiments, the trajectories may be assembled by the vehicle before transmitting them to server 1230. In other embodiments, the trajectories may be assembled server side, for example, based on data received from vehicles traversing road segment 3600. For example, in some embodiments, the vehicles (e.g., vehicles 200) may transmit data to one or more servers relating to their motion along road segment 3600 (e.g., steering angle, heading, time, position, speed, sensed road geometry, and/or sensed landmarks, among other things). The server may reconstruct trajectories for the vehicles based on the received data.

[0473] In some embodiments, the lane split may be identified based on a first trajectory, such as trajectory 3701, and a second trajectory, such as trajectory 3704. In such embodiments, the lane split may be identified based on a lateral spacing between trajectories 3701 and 3704. For example, the lane split may be detected based on an increase in the lateral spacing between the trajectories after the lane split. In some embodiments, the lane split may be determined by comparing the lateral spacing between trajectories 3701 and 3704 to a lateral spacing threshold. For example, the lane split may be identified based on the lateral spacing exceeding a threshold (e.g., 0.55m, 1.0m, 2.0m, etc.). In some embodiments, a rolling average, a regression model (linear regression, local regression, etc.), or various other statistical analysis methods may also be used to more accurately identify the lane split.

[0474] The lane split may also be determined based on environmental information or other data associated with road segment 3600. Server 1230 may use environmental data to verify whether the

divergence of trajectories represents a lane split rather than a lane change, etc. In some embodiments the environmental information may include one or more images captured by vehicles traversing road segment 3600. Server 1230 may analyze the images to identify representations of features of road segment 3600 indicative of the lane split. For example, server 1230 may determine that the shape of a lane boundary 3740 indicates that new lane 3610 has formed. As another example, server 1230 may detect dashed lane marks 3741 separating trajectories 3701 and 3704 after the divergence, but not before the divergence, which may indicate the lane split. Server 1230 may detect and/or verify the lane split based on other features in the images, such as a structure of the road (e.g., an increase in overall width, etc.), lane markings, vehicles, traffic signs, traffic arrows, traffic lights, etc. In some embodiments, the environmental information may not be an image but may include data representing any of the features above. These features may be detected, for example, using a trained model (e.g., using a deep neural network, etc.) to recognize features in captured images. In some embodiments, the environmental information may be received from vehicles traversing road segment 3600. For example, the environmental information may be received from the vehicles associated with trajectories 3701 and 3704, or may be received from other vehicles. In other embodiments, the environmental information may be retrieved from a storage location for determining the lane split. In other embodiments, the lane split feature may be determined without the use of environmental information (e.g., images, etc.) and may be based on analysis of the vehicle trajectories.

[0475] In some embodiments, the lane split may be determined based on multiple trajectories associated with vehicles entering lane 3610 and vehicles continuing along lane 3611. For example, a lane split feature may be detected based on crowdsourcing trajectories of vehicles traversing the road segment. Server 1230 may require sufficient evidence or data indicating a road split feature before the feature is mapped. For example, a single vehicle or only a few vehicles having diverging trajectories may be outliers and may not be sufficient to establish a lane split feature. Accordingly, server 1230 may perform a clustering or other operation for identifying the lane split based on the trajectories. FIG. 37B illustrates example clustering of vehicle trajectories for determining a lane split, consistent with the disclosed embodiments. As discussed above, trajectories 3701, 3702, and 3703 may be associated with vehicles entering new lane 3610 after the lane split, and trajectories 3704, 3705, and 3706 may be associated with vehicles continuing to travel within lane 3611. While three trajectories for each option (e.g., entering lane 3610 or continuing in lane 3611) are shown in FIG. 37B for illustration purposes, it is understood that any suitable number of trajectories may be used to identify the lane split.

[0476] Server 1230 may be configured to cluster trajectories 3701, 3702, 3703, 3704, 3705, and 3706 when detecting the lane split feature. For example, server 1230 may determine clusters 3721, 3722, 3723, 3724, 3725, and 3726. The clustering may be determined at various locations along road segment 3600. In some embodiments, the clustering may be determined at predetermined intervals (e.g., 1 m, 2 m, 5 m, 10 m, 20 m, 50 m, etc.). Various factors may be used for determining the spacing. For example, shorter spacing between intervals may result in a higher computational load, especially as the number of trajectories being analyzed (or the amount of data associated with each trajectory, etc.)

increases. Greater spacing between intervals, however, may provide less accuracy, so the complexity of the road segment or of the trajectory paths may also be considered. In some embodiments, the clustering may not be determined at predefined intervals. For example, the locations may be random or based on characteristics of the road segment and/or the trajectories.

[0477] At each location, server 1230 may be configured to cluster the trajectories based on a lateral spacing between the trajectories. The clustering may be determined by defining a segment normal to the trajectories or to the road segment at a given location and identifying all trajectories that intersect the segment. For example, in cluster 3721, all of trajectories 3701, 3702, 3703, 3704, 3705, and 3706 may be clustered together, as shown. In some embodiments, the clustering may be based on a minimum or threshold distance between the trajectories (e.g., 0.1 m., 0.5 m., 1.0 m., 2.0 m., 3.0 m., etc.). For example, in cluster 3722, all of the trajectories may again be clustered together, despite some divergence beginning to appear between trajectories 3701, 3702, and 3703 and trajectories 3704, 3705, and 3706. At clusters 3723 and 3725, however, the two groups of trajectories ((3701, 3702, and 3703; and 3704, 3705, and 3706) may be clustered separately, as shown. The same may be true for clusters 3724 and 3726. Based on the clustering, server 1230 may determine a divergence between the trajectories beginning at cluster 3722. This divergence may be indicative of the lane split. As discussed above, the lane split may be determined based on a lateral spacing between the clusters. For example, the lane split may be identified based on the lateral spacing exceeding a threshold (e.g., 0.5 m., 1.0 m., 2.0 m., etc.). In some embodiments, a rolling average, other statistical analysis methods may also be used to identify the lane split based on the clusters. Further, as discussed above, environmental information (e.g., images, etc.) may be used to confirm the divergence identified based on clustering. For example, server 1230 may use images to determine whether a divergence is associated with a lane split or if it should be considered an outlier, etc.

[0478] Server 1230 may further determine target trajectories corresponding to a drivable path to be used by autonomous vehicles through road segment 3600. In some embodiments, this may include connecting the series of clusters developed based on the trajectories. FIG. 37C illustrates example target trajectories 3730 and 3731 consistent with the disclosed embodiments. Target trajectory 3730 may be associated with new lane 3601, whereas target trajectory 3731 may be associated with lane 3611. The target trajectories may be represented as a spline with nodes at each of the cluster locations, as shown in FIG. 37C. In some embodiments, the trajectories may be represented in other forms, such as a series of points, etc. The points may be derived from splines either by server 1230 or locally in the vehicles. In some embodiments, the target trajectories may be determined based on the number of trajectories connecting each of the clusters. For example, trajectories 3701, 3702, and 3703 may include both cluster 3722 and 3723 (see FIG. 37B), and thus server 1230 may form target trajectory 3730 to include clusters 3722 and 3723. In some embodiments, server 1230 may include a connection between pairs of clusters based on a minimum number of trajectories connecting the two clusters. As a result, target trajectories 3730 and 3731 may be determined. Target trajectory 3730 may be begin at a branch point of divergence point corresponding with cluster 3722.

[0479] In some embodiments, server 1230 may further be configured to adjust various points along trajectories 3730 and 3731. For example, because trajectories 3730 and 3731 are developed based on the clustering technique described above, the point at which trajectories 3730 and 3731 begin to diverge will be at the last point in which the two groups of trajectories are clustered together. In reality, however, the actual trajectories of the vehicles entering new lane 3610 may occur earlier in the trajectory (e.g., between cluster 3722 and 3721), or later in the trajectory (e.g., in the case of a lane merge). Accordingly, the resulting trajectory 3730 may branch off from trajectory 3731 at a more abrupt angle than a vehicle would normally drive. To correct or account for this, server 1230 may adjust the location of the divergence.

[0480] FIG. 38A illustrates an example process for adjusting the location of a branch point consistent with the disclosed embodiments. Adjustment of the branch point may be performed in a similar manner to the clustering technique discussed above. For example, server 1230 may examine all trajectories associated with a particular branch or divergence point. In the example shown in FIG. 38A, this may include trajectories 3701, 3702, 3703, 3704, 3705, and 3706. Server 1230 may sample the trajectories at a plurality of longitudinal distances before and/or after the cluster point defining the divergence (e.g., cluster 3722). The longitudinal distances may be spaced at a predetermined distance (e.g., 0.1 m., 0.2 m., 0.5 m., 1.0 m., etc.), which may be shorter than the predetermined distance for the clustering process discussed above. At each of the longitudinal distances, server 1230 may determine a lateral distance, D , between group 3801 of trajectories 3701, 3702, and 3703 and group 3802 of trajectories 3704, 3705, and 3706. Server 1230 may then adjust the divergence point based on where group 3801 and 3802 begin to diverge. FIG. 38B illustrates example adjusted target trajectories 3730 and 3731 consistent with the disclosed embodiments. For example, server 1230 may determine that target trajectory 3730 branches from target trajectory 3731 at branch point 3810. Various other points along target trajectories 3730 and 3731 may similarly be adjusted based on the actual trajectories. Various other means for adjusting the branch point may also be used. For example, in some embodiments, the branch point may be moved backwards by a predetermined factor or distance. In other embodiments, the branch point may be adjusted based on analysis of the cluster points alone, without further sampling. For example, the location of the branch point may be estimated based on the angle formed by clusters 3723, 3722, and 3725, etc. In other embodiments, the branch point may be adjusted by comparing the angle between target trajectories 3730 and 3731 to a threshold to ensure the branch angle is not too sharp, etc.

[0481] In some embodiments, determining target trajectories 3730 and 3731 may include correcting or removing one or more anomalies in the target trajectories. Figs. 39A-39C illustrate example anomalies that may arise in target trajectories 3730 and 3731, consistent with the disclosed embodiments. Such anomalies may arise due to inconsistencies in trajectories traveled by vehicles along road segment 3600. In some instances, the anomalies may include a divergence between target trajectory 3730 and target trajectory 3731 that is not indicative of the lane split feature. For example, FIG. 39A shows a segment of a target trajectory where the target trajectory diverges and converges, which may indicate that the divergence is not representative of a lane split. Such anomalies may be detected in various other

ways, for example, by analyzing the lane split feature in reference to other data (e.g., mapped lane marks, image data, mapped landmarks, etc.), through manual analysis, etc. The anomaly shown in FIG. 39A may be corrected, for example, by removing at least one of nodes 3901 and 3902 or combining the nodes to form a continuous target trajectory. Similarly, as shown in FIG. 39B, a branch point between two target trajectories may include anomalous nodes 3903 and/or 3904. The anomaly may be corrected, for example, by removing at least one of node 3903 or 3904 or combining the nodes. In other instances, the target trajectories may include leaves and/or gaps in the trajectories, as shown in FIG. 39C. Such anomalies may be corrected by connecting nodes 3905 and 3906 or, where the leave is not part of a lane split feature, removing extraneous node 3905. In some embodiments, additional nodes may be added as needed.

[0482] As a result of the steps described above with respect to Figs. 37A-39C, target trajectories may be determined representing drivable paths associated with the lane split feature of road segment 3600. It is understood that the various steps or techniques may not necessarily be performed in the order discussed above. For example, correcting anomalous points in the target trajectories may be performed prior to adjusting the divergence point, etc. In some embodiments, server 1230 may update a vehicle road navigation model to include target trajectories 3730 and 3731. Server 1230 may transmit the vehicle road navigation model or the update to the model to the vehicles that transmitted data to server 1230. Server 1230 may transmit the autonomous vehicle road navigation model or the update to the model to other vehicles that travel on road segment 3600 at later times. In some embodiments, server 1230 may also update and/or distribute a sparse map, such as sparse map 800 discussed above, based on the target trajectories.

[0483] While a lane split feature is used by way of example throughout the description above, it is to be understood that the disclosed embodiments may also apply to lane merges. For example, using the clustering and other techniques above, server 1230 may be configured to detect vehicle trajectories merging together along a road segment. For example, a lane merge feature may include an entrance ramp where vehicles merge onto a freeway, a road narrowing, the end of a turn or merge lane, a "Y" junction in a road, or various other lane merge features. One skilled in the art would appreciate various modifications to the techniques discussed above may be applicable for detecting lane merges rather than splits. For example, the system may be configured to detect a convergence between trajectories based on a lateral spacing falling below a threshold distance, a branch point may be adjusted forwards, rather than backwards along an intended direction of travel of a road segment, etc.

[0484] In some implementations, the disclosed embodiments may be used to detect an obstacle or blockage in a road segment, such as a stalled vehicle, an animal, an emergency or service vehicle, a road construction project, a lane closure, etc. Based on the trajectories of vehicles traversing the obstacle in the road, server 1230 may detect a lane split and/or merge associated with the obstacle. In some embodiments, server 1230 may be configured to distinguish lane split features from road obstacles, etc. For example, server 1230 may detect the road obstacle based on the abruptness of the divergence in vehicle trajectories, the convergence of the trajectories after the split, the timing of the divergence in

vehicle trajectories, compared to historical data (e.g., if the detected lane split appears abruptly, etc.).

Various control actions may be performed based on a detected obstacle. For example, server 1230 may not update the road navigation model or may update it temporarily. Server 1230 may also transmit a notification of the obstacle, include an indicator of the obstacle in the road navigation model, etc.

[0485] The disclosed embodiments may include systems and methods for mapping lane merges, as disclosed herein. For example, a system for mapping lane merges for use in vehicle navigation may include at least one processor programmed to receive first navigational information from a first vehicle that has navigated along a road segment. The road segment may include a lane merge feature, and the road segment may include at least a first travel lane and a second travel lane prior to the lane merge feature that transition into a third travel lane after the merge feature. The at least one processor may be further programmed to receive second navigation information from a second vehicle that has navigated along the road segment and to receive at least one image associated with the road segment. The at least one processor may be further programmed to determine, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the third travel lane of the road segment, to determine, from the second navigational information, a second actual trajectory of the second vehicle along the second travel lane and the third travel lane of the road segment, determine a convergence between the first actual trajectory and the second actual trajectory, and determine, based on analysis of the at least one image, that the convergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane merge feature. The at least one processor may be further programmed to update a vehicle road navigation model to include a first target trajectory corresponding to the first travel lane prior to the lane merge feature and extending along the third travel lane after the lane merge feature, and a second target trajectory that extends along the second travel lane before the lane merge feature and joins with the first target trajectory.

[0486] The disclosed embodiments may also include systems and methods for navigating a lane merge feature. For example, a system for navigating a host vehicle along a road segment may include at least one processor programmed to receive, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory corresponding to a first travel lane prior to a lane merge feature and extending along a third travel lane after the lane merge feature, and wherein the vehicle road navigation model also includes a second target trajectory that extends along a second travel lane before the lane merge feature and joins with the first target trajectory. The at least one processor may be further programmed to receive information indicative of an environment of the host vehicle, determine, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory, and determine a navigational action to cause the host vehicle to navigate along the determined target trajectory.

[0487] Consistent with the disclosed embodiments, a host vehicle may be configured to navigate a road segment containing a lane split feature based on the updated road navigation model received from server 1230. FIG. 40A illustrates an example road segment 4000 including a lane split feature, consistent with the disclosed embodiments. Host vehicle 4010, which may correspond to vehicle 200, as discussed

above, may be traversing road segment 4000 along target trajectory 4031, which may be represented in a road navigation model. The model may further include a second target trajectory 4031, representing a lane split feature. Target trajectory 4030 and/or target trajectory 403 may have been developed using the lane split mapping techniques discussed above.

5 [0488] As host vehicle 4010 approaches the lane split feature, it may determine whether to continue along target trajectory 4031 or to follow target trajectory 4030. Host vehicle 4010 may receive or access information in order to make this determination. For example, such information may be indicative of an environment of the vehicle. In some embodiments, the information may be accessed by a server, such as server 1230, which may provide information about a route of the host vehicle, properties of road segment 4000, properties of target trajectories 4030 and 4031, or various other information. In 10 some embodiments, the information may be captured by the vehicle, for example, using an image capture device, one or more sensors, information input to a user interface, etc. The information may include properties of the vehicle (e.g., a current speed, a current position, a heading direction, a fuel level of the vehicle, a maintenance condition of the vehicle, or various other operating parameters), or properties of 15 the environment (e.g., a speed limit, a road condition, a weather condition, the presence of target vehicles on the road segment, a speed or other properties of a target vehicle, detected landmarks, obstacles in the road, etc.). In some embodiments, the information may be stored locally on host vehicle 4010, for example, on memory 140, in map database 160, or the like, and may be accessed by vehicle 4010.

[0489] As shown in FIG. 40A, target trajectory 4030 may represent a widening of road segment 20 4000, such as the addition of a passing lane, etc. Host vehicle 4010 may determine whether to navigate along target trajectory 4030 based on information indicative of the environment of host vehicle 4010. For example, host vehicle 4010 may be traveling behind target vehicle 4021 and may determine whether to navigate along target trajectory 4030 based on information about target vehicle 4021. In some embodiments, host vehicle 4010 may determine a speed of target vehicle 4021 and may determine 25 whether to pass vehicle based on the determined speed. For example, if target vehicle 4021 is traveling slower than host vehicle 4010 (or slower than a target or desired speed of host vehicle 4010), host vehicle 4010 may determine to travel along target trajectory 4030 in order to pass target vehicle 4021. In some embodiments, the determination may also be based on one or more rules. For example, host vehicle 4010 may pass target vehicle 4021 only if doing so would not cause host vehicle 4010 to exceed a speed limit, 30 a maximum acceleration, etc.

[0490] In some embodiments, as discussed above, host vehicle 4010 may determine the 3 information based on images captured by host vehicle 4010. FIG. 40B illustrates an example image 4020 that may be used by a host vehicle in determining a target trajectory associated with a lane split feature, consistent with the disclosed embodiments. Host vehicle may detect target vehicle 4021 based on the 35 image, as shown. Host vehicle may determine other properties of target vehicle 4021, such as a position, speed, acceleration, size, heading direction, etc., which may be used to determine whether to travel along target trajectory 4030 or target trajectory 4031. Host vehicle 4010 may further detect other features from

the image, such as lane mark 4022 or other landmarks, which may be used to determine properties of target vehicle 4021 or to determine which target trajectory to follow.

[0491] In some embodiments, the information used by host vehicle 4010 to determine which target trajectory to follow may be based on an intended route or destination of host vehicle 4010. For example, host vehicle 4010 may determine whether target trajectory 4030 is consistent with the target route. FIG. 40C illustrates another example road segment 4001 including a lane split feature, consistent with the disclosed embodiments. In this example, the lane split feature may be associated with an exit ramp. Target trajectory 4030 may be associated with exiting the current road or highway and host vehicle 4010 may determine whether target trajectory 4030 is consistent with the intended route. For example, taking the exit ramp may provide a faster or more direct route to the target destination of the host vehicle. While an example of an exit ramp is used in FIG. 40C, the road split feature may be a turn lane, a "Y" junction in the road, or various other lane split features, as discussed above. As with the previous example, host vehicle 4010 may determine whether to follow target trajectory 4030 based on information indicative of the environment of host vehicle 4010. This may include images captured by an image capture device of host vehicle 4010. FIG. 40D illustrates another example image 4040 that may be used by a host vehicle in determining a target trajectory associated with a lane split feature, consistent with the disclosed embodiments. In this example, host vehicle may detect a landmark, such as exit sign 4041, which may indicate that target trajectory 4030 is associated with an exit lane. Various other features of road segment 4001 may be detected, including a road structure (e.g., the width of the road, whether the lanes diverge, etc.), lane markings, weather conditions, road conditions, obstacles, target vehicles, etc. In some embodiments, host vehicle 4010 may analyze image 4040 to determine whether there is traffic or other conditions along target trajectory 4030 or 4031, which may inform the determination of which target trajectory to follow. It is to be understood that the example information discussed above is provided by way of example, and other forms or information may be used in determining how to navigate the lane split feature.

[0492] Based on the determined target trajectory, host vehicle 4010 may determine a navigation action consistent with the target trajectory. For example, if host vehicle 4010 determines to follow target trajectory 4030, the navigation action may be a steering action to guide vehicle 4010 along target trajectory 4030. Where vehicle 4010 determines that it should follow target trajectory 4031, the navigation action may be to maintain the current heading direction. Various other navigation actions may be included, such as a braking action, an acceleration, a deceleration, a merge maneuver, etc.

[0493] FIG. 41 is a flowchart depicting an exemplary process 4100 for mapping lane splits for use in vehicle navigation, consistent with the disclosed embodiments. Process 4100 may be performed by at least one processing device included on a server, such as server 1230. Process 4100 is exemplary only. One of ordinary skill in the art having the benefit of this disclosure may understand that process 4100 may include additional steps, exclude certain steps, or may be otherwise modified in a manner consistent with this disclosure.

[0494] At step 4110, process 4100 may include receiving first navigational information from a first vehicle that has navigated along a road segment, wherein the road segment includes a lane split feature. The road segment may include at least a first travel lane prior to the lane split feature that transitions into at least a second travel lane and a third travel lane after the lane split feature. For example, the road segment may correspond to road segment 33600 and thus the second lane of travel may correspond to new lane 33610 and the third lane of travel may correspond to lane 33611. In some embodiments, the lane split feature may correspond to a widening of the road segment. Accordingly, the second travel lane and the third travel lane may extend parallel to one another after the lane split feature. Accordingly, the second travel lane and third travel lane may continue to be part of the same road segment. In other embodiments, the lane split feature may correspond to a turn lane, a "Y" junction, an exit ramp, or the like. Accordingly, the second travel lane and the third travel lane may diverge from one another after the lane split feature. For example, the second travel lane may be associated with a separate road segment after the split. In some embodiments, the first travel lane and the third travel lane may form a continuous lane of travel. As discussed, above, the first navigational information may be uploaded or transmitted to server 1230 by the first vehicle. The first navigational information may comprise location identifiers recorded by the first vehicle while traversing the road segment.

[0495] At step 4111, process 4100 may include receiving second navigation information from a second vehicle that has navigated along the road segment. Similar to the first navigation information, the second navigation information may be uploaded or transmitted to server 1230 by the second vehicle. The second navigational information may comprise location identifiers recorded by the second vehicle while traversing the road segment. In some embodiments the first vehicle and second vehicle may be the same vehicle, for example, traversing the road segment at different times. Alternatively, the first vehicle and second vehicle may be different vehicles.

[0496] At step 4112, process 4100 may include receiving at least one image associated with the road segment. In some embodiments, the image may be captured by a host vehicle traversing the road segment. For example, the at least one image may be captured by an image capture device of the first vehicle or the second vehicle. The at least one image may depict features of the environment of the host vehicle and may include a representation of the lane split feature or other associated features. In some embodiments, the at least one image may be retrieved from a server, for example, from a storage device of server 1230.

[0497] At step 4113, process 4100 may include determining, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the second travel lane of the road segment. For example, step 4113 may comprise analyzing the location identifiers recorded by the first vehicle to determine the first actual trajectory. The first actual trajectory may be represented by a spline having a plurality of nodes, by a plurality of points, or in various other formats.

[0498] At step 4114, process 4100 may include determining, from the second navigational information, a second actual trajectory of the second vehicle along the first travel lane and the third travel lane of the road segment. Similar to step 4113, step 4114 may comprise analyzing the location identifiers

recorded by the second vehicle to determine the second actual trajectory. The second actual trajectory may be represented by a spline having a plurality of nodes, by a plurality of points, or in various other formats.

[0499] At step 415, process 4100 may include determining a divergence between the first actual trajectory and the second actual trajectory. As discussed above, the divergence may be determined based on a lateral spacing between the first actual trajectory and the second actual trajectory. Further, the divergence may be determined based on a comparison of the lateral spacing to a threshold.

[0500] At step 416, process 4100 may include determining, based on an analysis of the at least one image, that the divergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane split feature. In some embodiments, the determination that the divergence between the first actual trajectory and the second actual trajectory is indicative of the presence in the road segment of the lane split feature is further based on a determined lateral spacing between the first actual trajectory and the second actual trajectory. For example, server 1230 may analyze the lateral spacing between the first actual trajectory and the second actual trajectory to determine whether the divergence is significant enough to correspond to a lane split feature. Server 1230 may also analyze the at least one image to determine if the divergence is indicative of the lane split feature. For example, determining that the divergence between the first actual trajectory and the second actual trajectory is indicative of the presence in the road segment of the lane split feature may be further based on a representation of a lane mark, a structure of the road, or other features detected in the at least one image.

[0501] In some embodiments, process 4100 may include analyzing multiple actual trajectories. Accordingly, process 4100 may include clustering and other techniques discussed above. For example, process 4100 may further include receiving a third actual trajectory that has navigated along the road segment. Process 4100 may further comprise clustering the third actual trajectory with at least one of the first actual trajectory or the second actual trajectory at a plurality of locations along the road segment. The first actual trajectory and the second actual trajectory may be clustered together in at least one first location before the lane split feature, and may be clustered separately in at least one second location after the lane split feature, similar to the clusters shown in FIG. 37B. As discussed above, the plurality of locations may be separated by a predetermined longitudinal distance, which may depend on computation power, complexity of the route, size of the vehicle navigation information, etc. Any other features or aspects of the clustering techniques discussed above may also apply to process 4100.

[0502] At step 417, process 4100 may include updating a vehicle road navigation model to include a first target trajectory and a second target trajectory. The first target trajectory may correspond to the first travel lane prior to the road split feature and may extend along the second travel lane after the lane split feature. The second target trajectory may branch from the first target trajectory and may extend along the third travel lane after the lane split feature. For example, the second target trajectory may begin at a branch point adjacent to or contiguous to the first target trajectory. In embodiments where clustering techniques are used, the first target trajectory and the second target trajectory may be determined based on

the clustering, as discussed above, with respect to FIG. 337C. In some embodiments, process 4100 may further comprise distributing the updated vehicle road navigation model to a plurality of vehicles, as discussed above.

[0503] Process 4100 may include various other techniques or steps as discussed above with respect to Figs. 337A-39C. For example, process 4100 may further include refining a position at which the second target trajectory branches from the first target trajectory by analyzing a distance between the first actual trajectory and the second actual trajectory at a plurality of locations surrounding the branch point to determine where the first actual trajectory and the second actual trajectory begin to diverge. Accordingly, the branch point may be adjusted to better correspond with the first actual trajectory and/or the second actual trajectory. In some embodiments, process 4100 may include refining a position at which the second target trajectory branches from the first target trajectory based on a third actual trajectory that has navigated along the road segment. For example, the third actual trajectory may provide additional information for adjusting the branch position.

[0504] In some embodiments, process 4100 may include removing at least one anomaly from at least one of the first target trajectory or the second target trajectory, as discussed above with respect to Figs. 39A-39C. For example, the at least one anomaly may include a divergence between the first target trajectory and the second target trajectory that is not indicative of the lane split feature. Removing the anomaly may include removing, adjusting, and/or combining one or more nodes associated with the first target trajectory and/or the second target trajectory.

[0505] FIG. 42 is a flowchart depicting an exemplary process 4200 for mapping lane splits for navigating a host vehicle along a road segment, consistent with the disclosed embodiments. Process 4200 may be performed by at least one processing device of a vehicle, such as processing unit 110, described above. Process 4200 is exemplary only. One of ordinary skill in the art having the benefit of this disclosure may understand that process 4200 may include additional steps, exclude certain steps, or may be otherwise modified in a manner consistent with this disclosure.

[0506] At step 4210, process 4200 may include receiving, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory that corresponds with a first lane of travel along the road segment prior to a lane split feature associated with the road segment and extends along a second lane of travel of the road segment subsequent to the lane split feature. The vehicle road navigation model may also include a second target trajectory that branches from the first target trajectory and extends along a third lane of travel of the road segment subsequent to the lane split feature. The vehicle road navigation model may correspond to the vehicle navigation model updated according to process 4200, as described above.

[0507] At step 4211, process 4200 may include receiving information indicative of an environment of the host vehicle. The information may be any form of information relevant to navigating the lane split feature, and may include any of the information discussed above with respect to Figs. 40A-40D. For example, the information may comprise at least one image captured by an image capture device of the host vehicle. The image capture device may correspond to image acquisition unit 1200

described above. In some embodiments, determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory may comprise identifying, based on the at least one image, a target vehicle to be passed by the host vehicle. For example, host vehicle 440 may determine whether to pass target vehicle 4402 as discussed above. Further, determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory may further comprise applying one or more rules or guidelines. The information may further comprise a speed of the target vehicle relative to the host vehicle, or various other properties of the target vehicle. In some embodiments, the information may comprise a planned route of the host vehicle. The information may be captured by a sensor or other device of the host vehicle, accessed from local storage associated with the host vehicle, or may be retrieved from a server, such as server 11230.

[0508] At step 4212, process 4200 may include determining, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory. In embodiments where the information comprises a planned route of the host vehicle, determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory may be based on whether the first target trajectory or the second target trajectory is consistent with the planned route. In some embodiments, determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory comprises determining that at least one of the first target trajectory or the second target trajectory is associated with a turn lane, an exit ramp, a passing lane, or the like.

[0509] At step 4213, process 4200 may include determining a navigational action to cause the host vehicle to navigate along the determined target trajectory. For example, where the determined target trajectory is along the second target trajectory, the navigational action may comprise a heading direction change of the host vehicle such that the host vehicle follows the second target trajectory following the lane split feature. In other embodiments, the navigational action may comprise maintaining a heading direction of the host vehicle such that the host vehicle follows the first target trajectory following the lane split feature. Various other navigational actions may also be performed, such as a braking maneuver, a passing maneuver, an acceleration, a merging maneuver, or the like.

[0510] 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 CDROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray, 4K Ultra HD Blu-ray, or other optical drive media.

[0511] 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, FHTML, FHTML/AJAX combinations, XML, or HTML with included Java applets.

- [0512] Moreover, while illustrative embodiments have been described herein, the scope of any and all embodiments having equivalent elements, modifications, commissions, 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 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.
- 10 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 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 mapping lane splits for use in vehicle navigation, the system comprising:

at least one processor programmed to:

receive first navigational information from a first vehicle that has navigated along a road
 5 segment, wherein the road segment includes a lane split feature, and wherein the road
 segment includes at least a first travel lane prior to the lane split feature that transitions
 into at least a second travel lane and a third travel lane after the lane split feature;
 receive second navigation information from a second vehicle that has navigated along the road
 segment;

10 receive at least one image associated with the road segment;

determine, from the first navigational information, a first actual trajectory of the first vehicle
 along the first travel lane and the second travel lane of the road segment;

determine, from the second navigational information, a second actual trajectory of the second
 vehicle along the first travel lane and the third travel lane of the road segment;

15 determine a divergence between the first actual trajectory and the second actual trajectory;

determine, based on analysis of the at least one image, that the divergence between the first actual
 trajectory and the second actual trajectory is indicative of a presence in the road segment
 of the lane split feature; and

update a vehicle road navigation model to include:

20 a first target trajectory corresponding to the first travel lane prior to the road split feature
 and extending along the second travel lane after the lane split feature; and

a second target trajectory that branches from the first target trajectory and extends along
 the third travel lane after the lane split feature.

2. The system of claim 1, wherein the determination that the divergence between the first actual trajectory
 25 and the second actual trajectory is indicative of the presence in the road segment of the lane split feature
 is further based on a determined lateral spacing between the first actual trajectory and the second actual
 trajectory.

3. The system of claim 1, wherein the second travel lane and the third travel lane extend parallel to one
 another after the lane split feature.

30 4. The system of claim 1, wherein the second travel lane and the third travel lane diverge from one
 another after the lane split feature.

5. The system of claim 1, wherein the first navigational information comprises location identifiers
 recorded by the first vehicle while traversing the road segment and the second navigational information
 comprises location identifiers recorded by the second vehicle while traversing the road segment.

35 6. The system of claim 1, wherein the at least one image is captured by an image capture device of the
 first vehicle or the second vehicle.

7. The system of claim 1, wherein the at least one image is retrieved from a server.

8. The system of claim 11, wherein the at least one processor is further programmed to refine a position at which the second target trajectory branches from the first target trajectory based on a third actual trajectory that has navigated along the road segment.

9. The system of claim 8, wherein the at least one processor is further configured to cluster the third actual trajectory with at least one of the first actual trajectory or the second actual trajectory at a plurality of locations along the road segment.

10. The system of claim 9, wherein the first actual trajectory and the second actual trajectory are clustered together in at least one first location before the lane split feature and are clustered separately in at least one second location after the lane split feature.

11. The system of claim 8, wherein the plurality of locations are separated by a predetermined longitudinal distance.

12. The system of claim 8, wherein the first target trajectory and the second target trajectory are determined based on the clustering.

13. The system of claim 1, wherein the at least one processor is further programmed to refine a position at which the second target trajectory branches from the first target trajectory by analyzing a distance between the first actual trajectory and the second actual trajectory at a plurality of locations surrounding the branch point to determine where the first actual trajectory and the second actual trajectory begin to diverge.

14. The system of claim 1, wherein the at least one processor is further programmed to distribute the updated vehicle road navigation model to a plurality of vehicles.

15. The system of claim 1, wherein the at least one processor is further programmed to remove at least one anomaly from at least one of the first target trajectory or the second target trajectory.

16. The system of claim 15, wherein the at least one anomaly includes a divergence between the first target trajectory and the second target trajectory that is not indicative of the lane split feature.

17. The system of claim 1, wherein determining that the divergence between the first actual trajectory and the second actual trajectory is indicative of the presence in the road segment of the lane split feature is further based on a representation of a lane mark detected in the at least one image.

18. The system of claim 1, wherein determining that the divergence between the first actual trajectory and the second actual trajectory is indicative of the presence in the road segment of the lane split feature is further based on a representation of a structure of the road detected in the at least one image.

19. A method for mapping lane splits for use in vehicle navigation, the method comprising:

receiving first navigational information from a first vehicle that has navigated along a road segment, wherein the road segment includes a lane split feature, and wherein the road segment includes at least a first travel lane prior to the lane split feature that transitions into, at least a second travel lane and a third travel lane after the lane split feature;

receiving second navigational information from a second vehicle that has navigated along the road segment;

receiving at least one image associated with the road segment;

determining, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the second travel lane of the road segment;

determining, from the second navigational information, a second actual trajectory of the second vehicle along the first travel lane and the third travel lane of the road segment;

5 determining a divergence between the first actual trajectory and the second actual trajectory;

determining, based on analysis of the at least one image, that the divergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane split feature; and

updating a vehicle road navigation model to include:

10 a first target trajectory corresponding to the first travel lane prior to the road split feature and extending along the second travel lane after the lane split feature; and
a second target trajectory that branches from the first target trajectory and extends along the third travel lane after the lane split feature.

20. A system for navigating a host vehicle along a road segment, the system comprising:

15 at least one processor programmed to:

receive, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory that corresponds with a first lane of travel along the road segment prior to a lane split feature associated with the road segment and extends along a second lane of travel of the road segment subsequent to the lane split feature, and wherein the vehicle road navigation model also includes a second target trajectory that branches from the first target trajectory and extends along a third lane of travel of the road segment subsequent to the lane split feature;

20

receive information indicative of an environment of the host vehicle;

determine, based on the information indicative of the environment of the host vehicle, whether to

25

navigate the host vehicle along the first target trajectory or the second target trajectory;

and

determine a navigational action to cause the host vehicle to navigate along the determined target trajectory..

30

21. The system of claim 20, wherein the information comprises at least one image captured by an image capture device of the host vehicle..

22. The system of claim 21, wherein determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory comprises identifying, based on the at least one image, a target vehicle to be passed by the host vehicle..

35

23. The system of claim 22, wherein the information further comprises a speed of the target vehicle relative to the host vehicle..

24. The system of claim 20, wherein determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory comprises determining that at least one of the first target trajectory or the second target trajectory is associated with a turn lane..

225. The system of claim 20, wherein determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory comprises determining that at least one of the first target trajectory or the second target trajectory is associated with an exit ramp.

226. The system of claim 20, wherein the information comprises a planned route of the host vehicle.

5 227. The system of claim 26, wherein determining whether to navigate the host vehicle along the first target trajectory or the second target trajectory is based on whether the first target trajectory or the second target trajectory is consistent with the planned route.

228. The system of claim 20, wherein the navigational action comprises maintaining a heading direction of the host vehicle such that the host vehicle follows the first target trajectory following the lane split feature.

229. The system of claim 20, wherein the navigational action comprises a heading direction change of the host vehicle such that the host vehicle follows the second target trajectory following the lane split feature.

30. A method for navigating a host vehicle along a road segment, the method comprising:

receiving, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory that corresponds with a first lane of travel along the road segment prior to a lane split feature associated with the road segment and extends along a second lane of travel of the road segment subsequent to the lane split feature, and wherein the vehicle road navigation model also includes a second target trajectory that branches from the first target trajectory and extends along a third lane of travel of the road segment subsequent to the lane split feature;

receiving information indicative of an environment of the host vehicle;

determining, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory; and

determining a navigational action to cause the host vehicle to navigate along the determined target trajectory.

31. A system for mapping lane merges for use in vehicle navigation, the system comprising: at least one processor programmed to:

receive first navigational information from a first vehicle that has navigated along a road segment, wherein the road segment includes a lane merge feature, and wherein the road segment includes at least a first travel lane and a second travel lane prior to the lane merge feature, that transition into a third travel lane after the merge feature;

receive second navigation information from a second vehicle that has navigated along the road segment;

receive at least one image associated with the road segment;

determine, from the first navigational information, a first actual trajectory of the first vehicle along the first travel lane and the third travel lane of the road segment;

determine, from the second navigational information, a second actual trajectory of the second vehicle along the second travel lane and the third travel lane of the road segment;

determine a convergence between the first actual trajectory and the second actual trajectory;

determine, based on analysis of the at least one image, that the convergence between the first actual trajectory and the second actual trajectory is indicative of a presence in the road segment of the lane merge feature; and

update a vehicle road navigation model to include:

a first target trajectory corresponding to the first travel lane prior to the lane merge feature and extending along the third travel lane after the lane merge feature; and

a second target trajectory that extends along the second travel lane before the lane merge feature and joins with the first target trajectory.

32. A system for navigating a host vehicle along a road segment, the system comprising:


at least one processor programmed to:

receive, from a server-based system, a vehicle road navigation model, wherein the vehicle road navigation model includes a first target trajectory corresponding to a first travel lane prior to a lane merge feature and extending along a third travel lane after the lane merge feature, and wherein the vehicle road navigation model also includes a second target trajectory that extends along a second travel lane before the lane merge feature and joins with the first target trajectory;

receive information indicative of an environment of the host vehicle;

determine, based on the information indicative of the environment of the host vehicle, whether to navigate the host vehicle along the first target trajectory or the second target trajectory; and

determine a navigational action to cause the host vehicle to navigate along the determined target trajectory.

**FIG. 1**




FIG. 2A




FIG. 2B





FIG. 2C




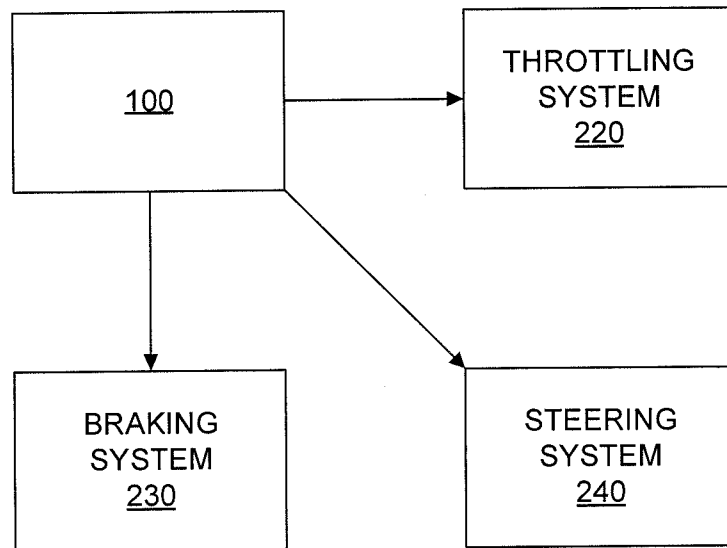


FIG. 2E

**FIG. 2F**

8/68




FIG. 3A





FIG. 3B

**FIG. 3C**



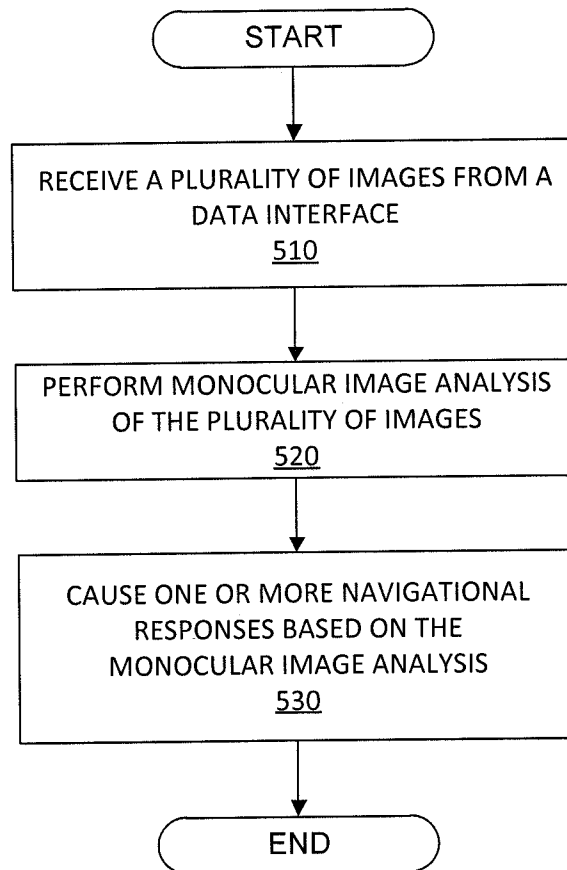
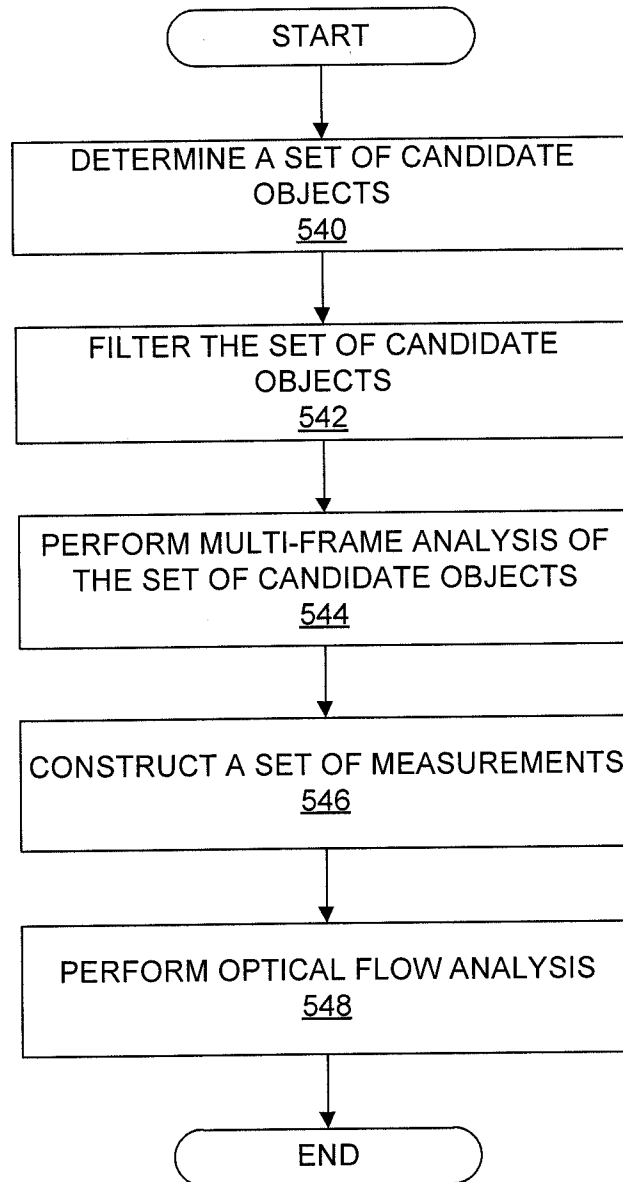


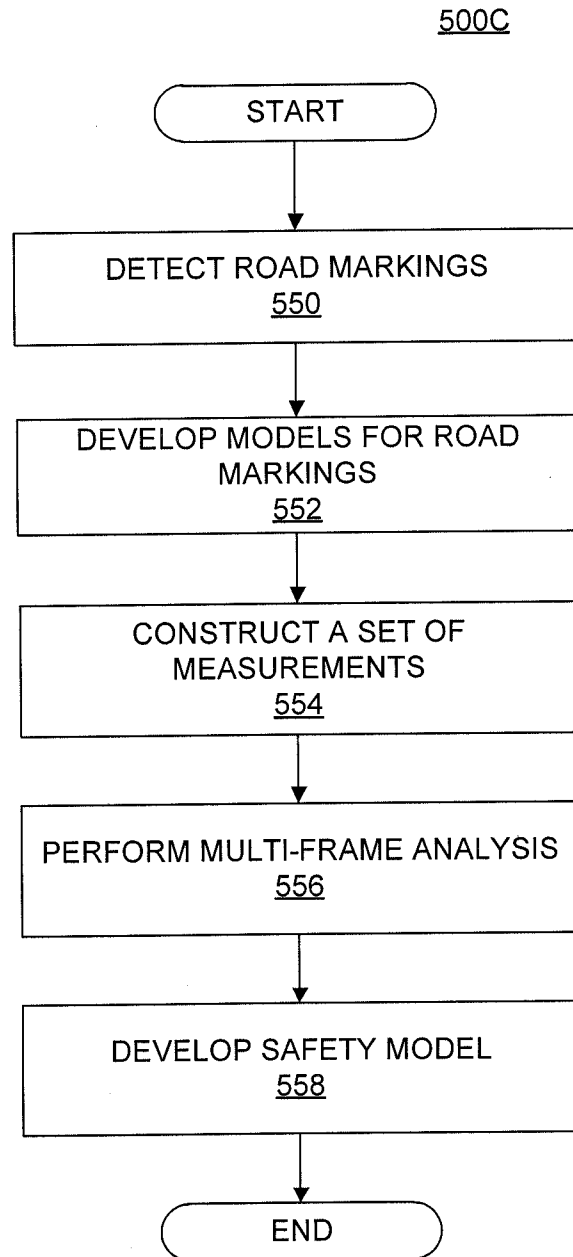
FIG. 3D

**FIG. 4**

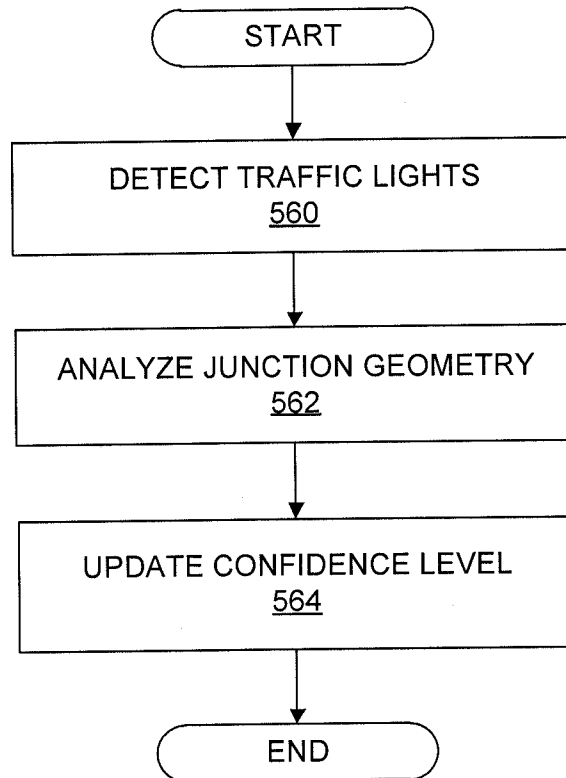
13/68

500A**FIG. 5A**

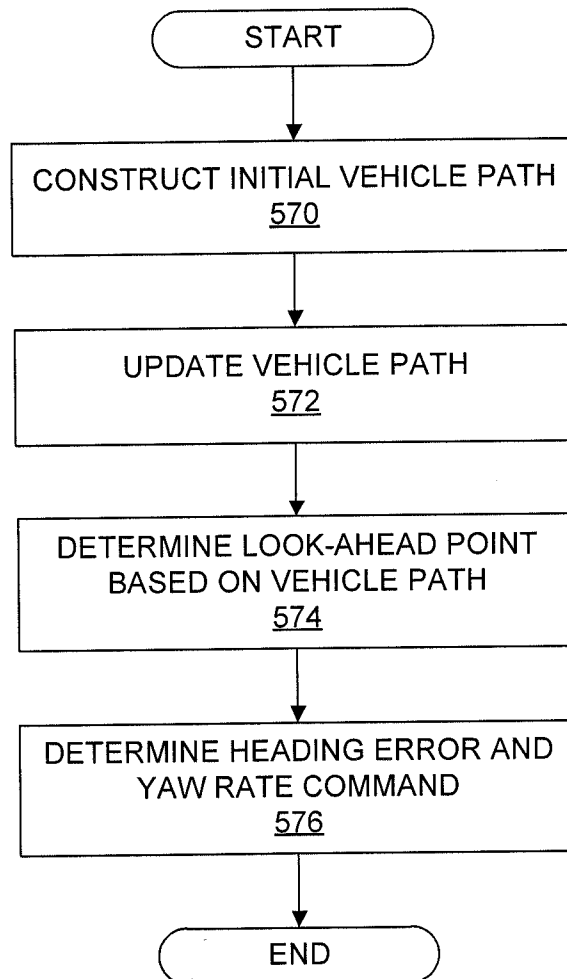
500B**FIG. 5B**

**FIG. 5C**

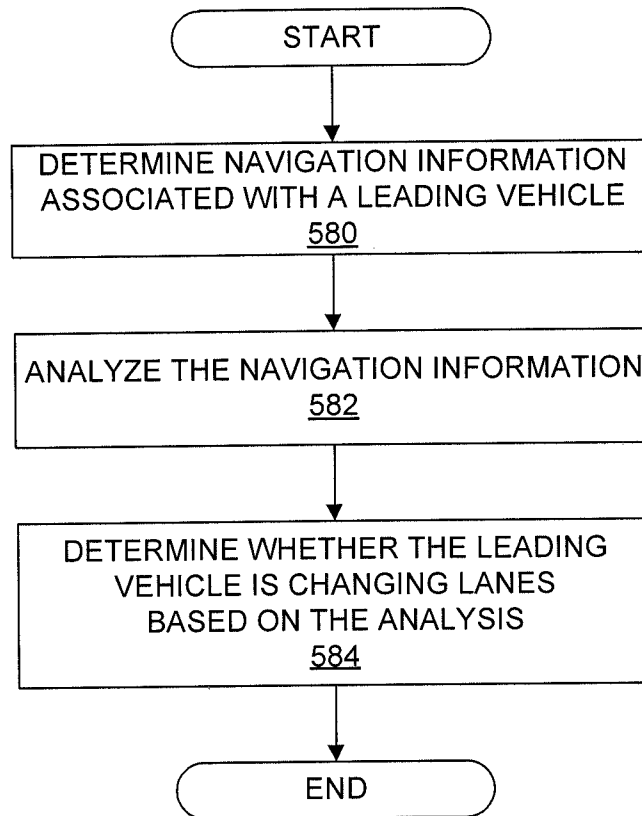
16/68

500D**FIG. 5D**

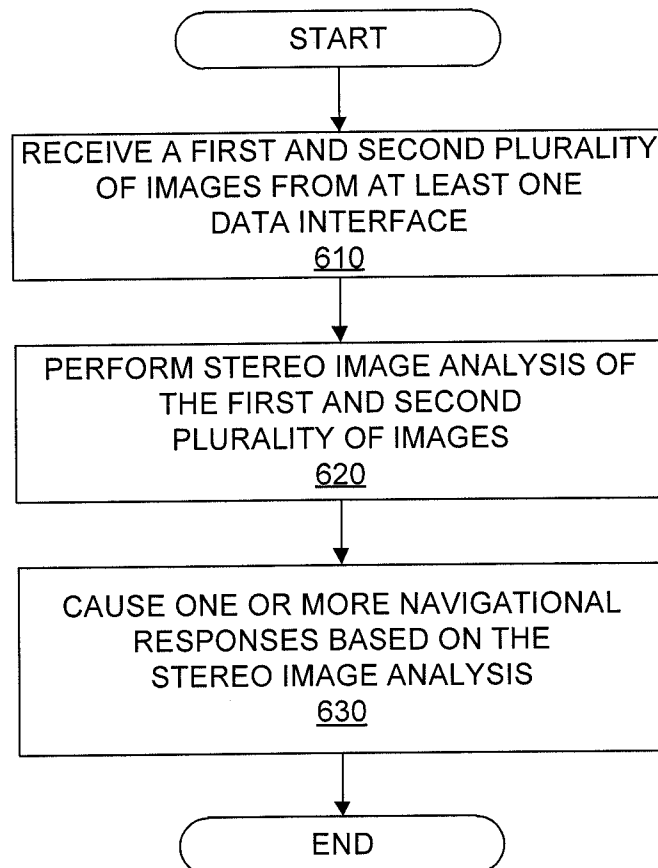
17/68

500E**FIG. 5E**

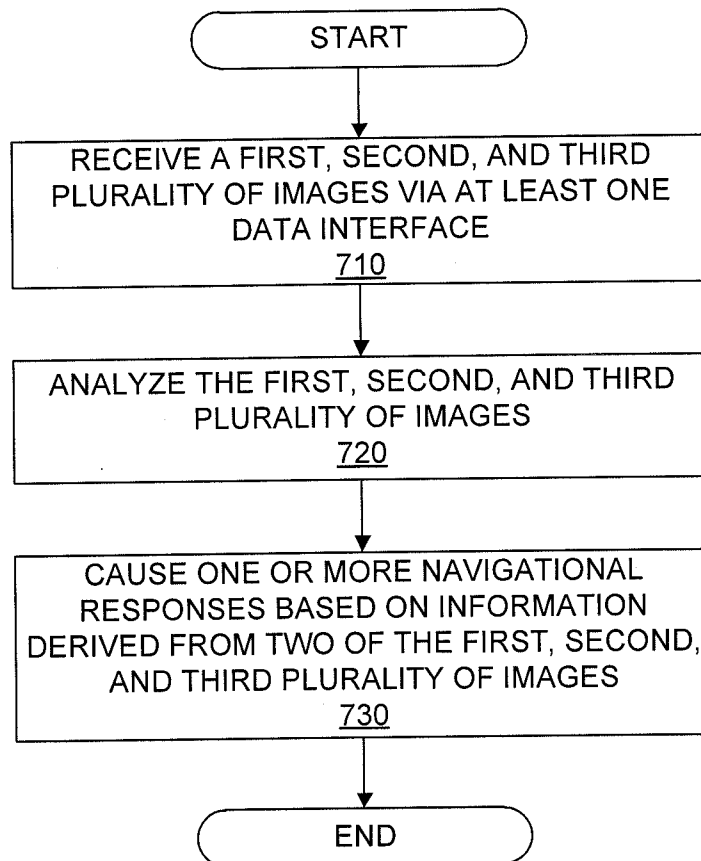
18/68

500F**FIG. 5F**

19/68

600**FIG. 6**

20/68

700**FIG. 7**




FIG. 8




FIG. 9A




FIG. 9B




FIG. 10




FIG. 11A




FIG. 11B




FIG. 11C




FIG. 11D




FIG. 12




FIG. 13




FIG. 14

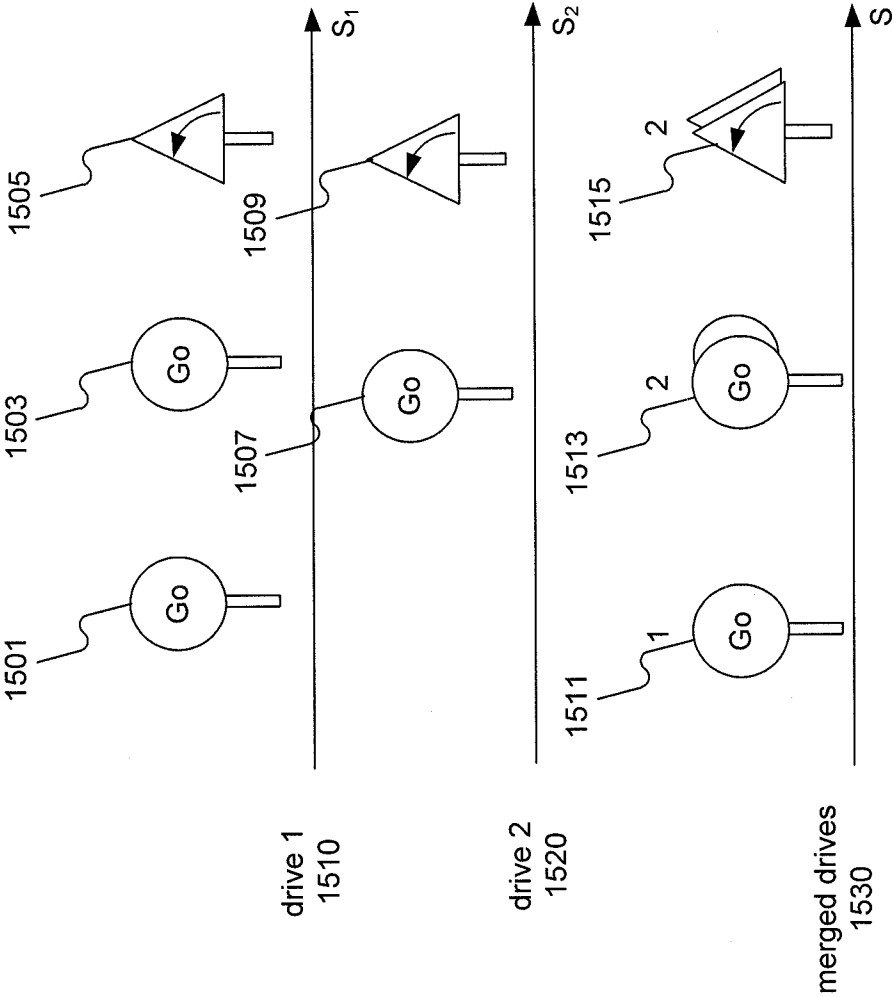


FIG. 15

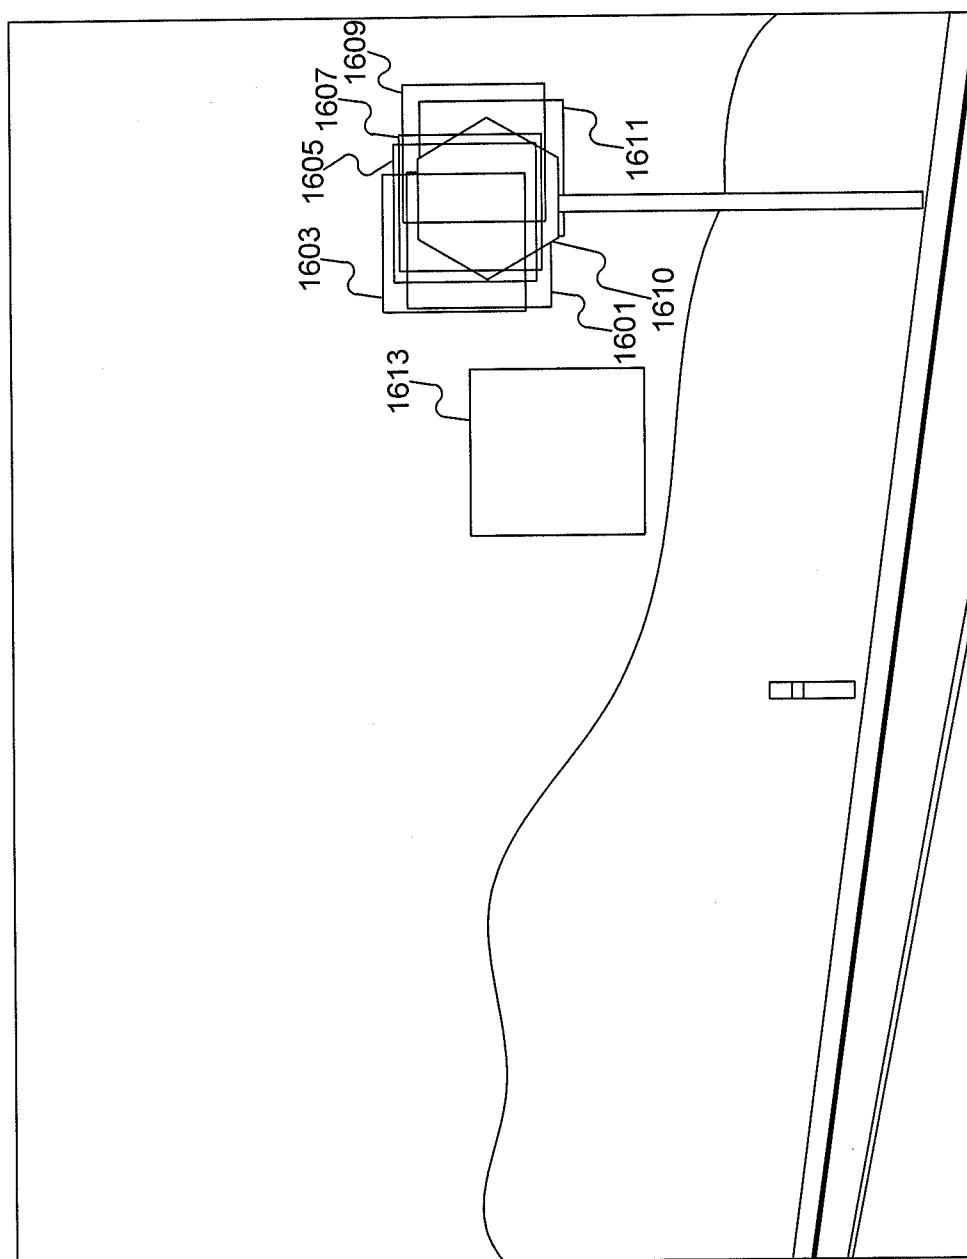


FIG. 16




FIG. 17




FIG. 18

35/68

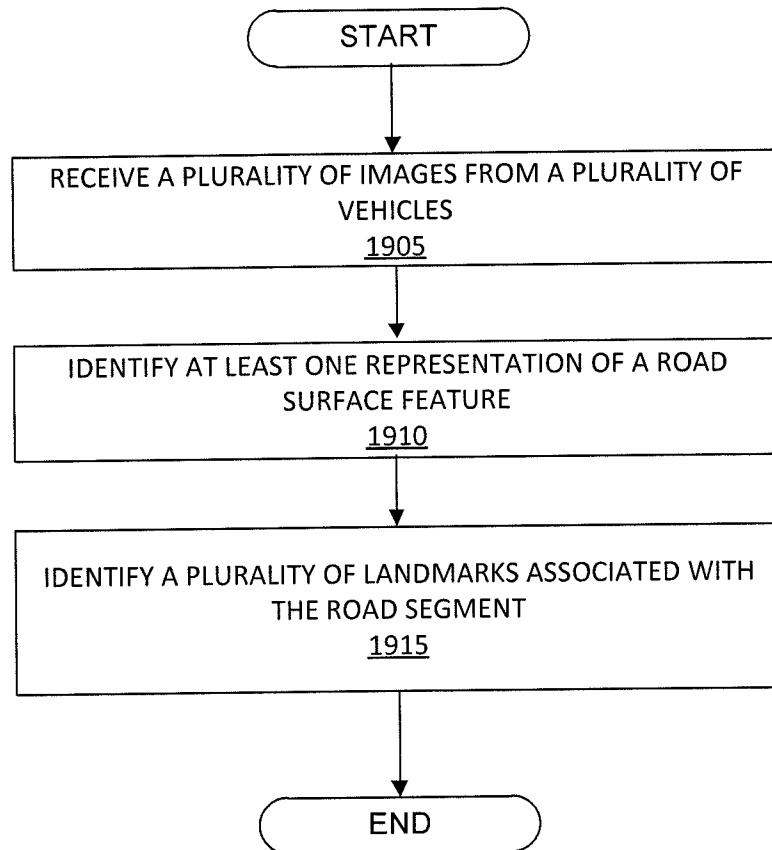
1900

FIG. 19

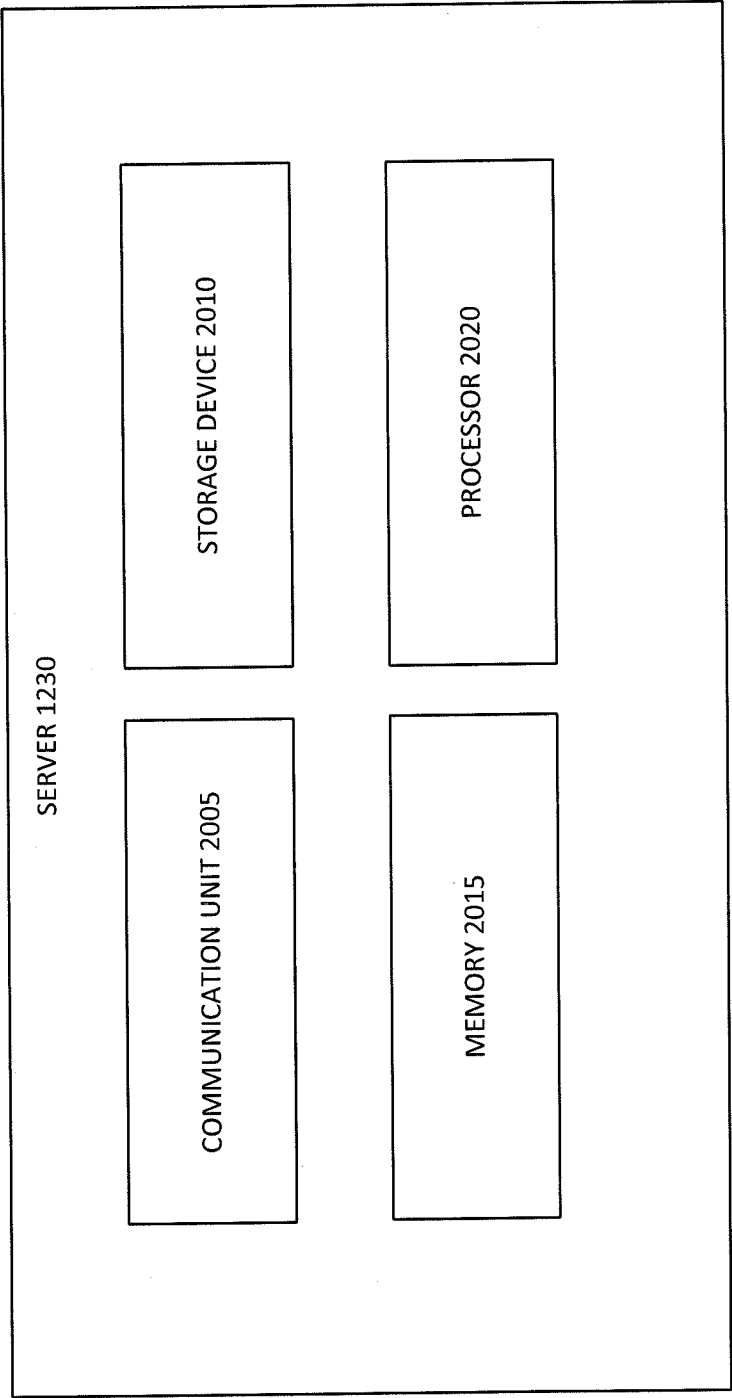


FIG. 20




FIG. 21




FIG. 22


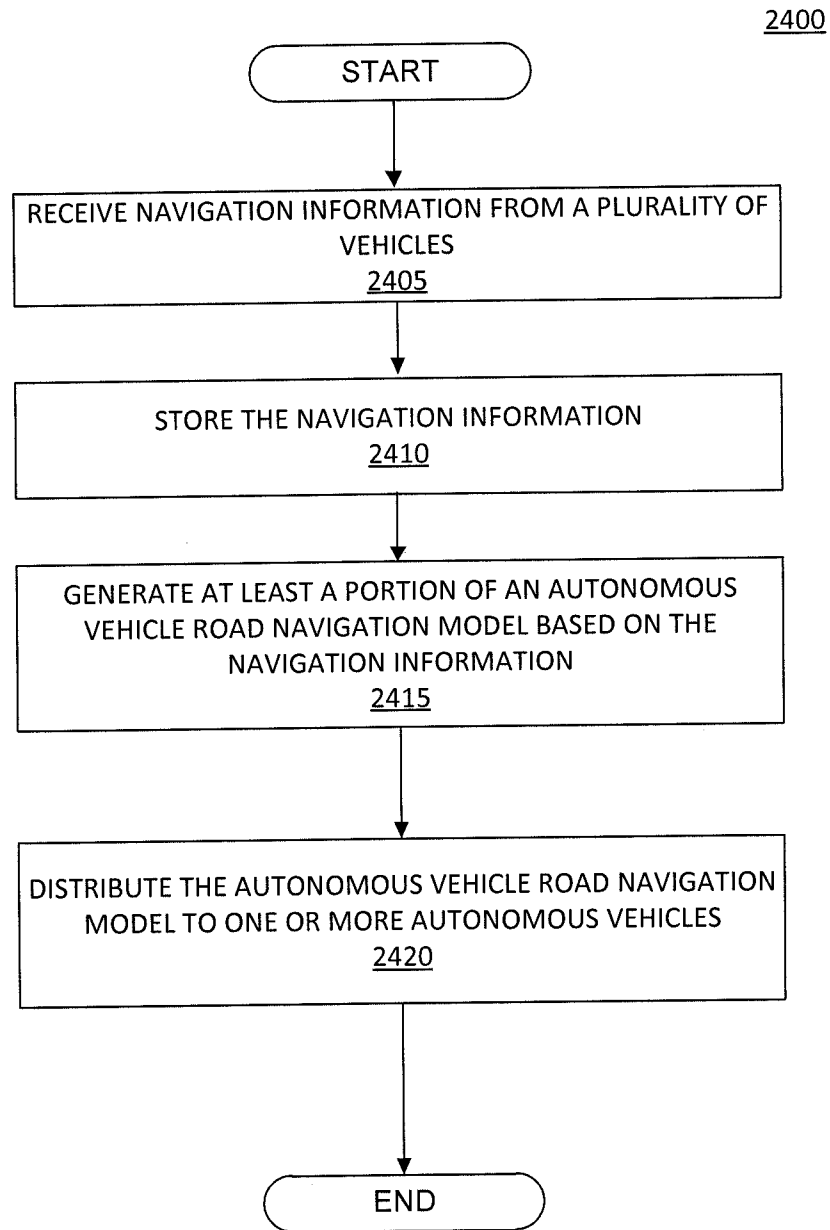



FIG. 23

40/68

**FIG. 24**

**FIG. 25**

42/68

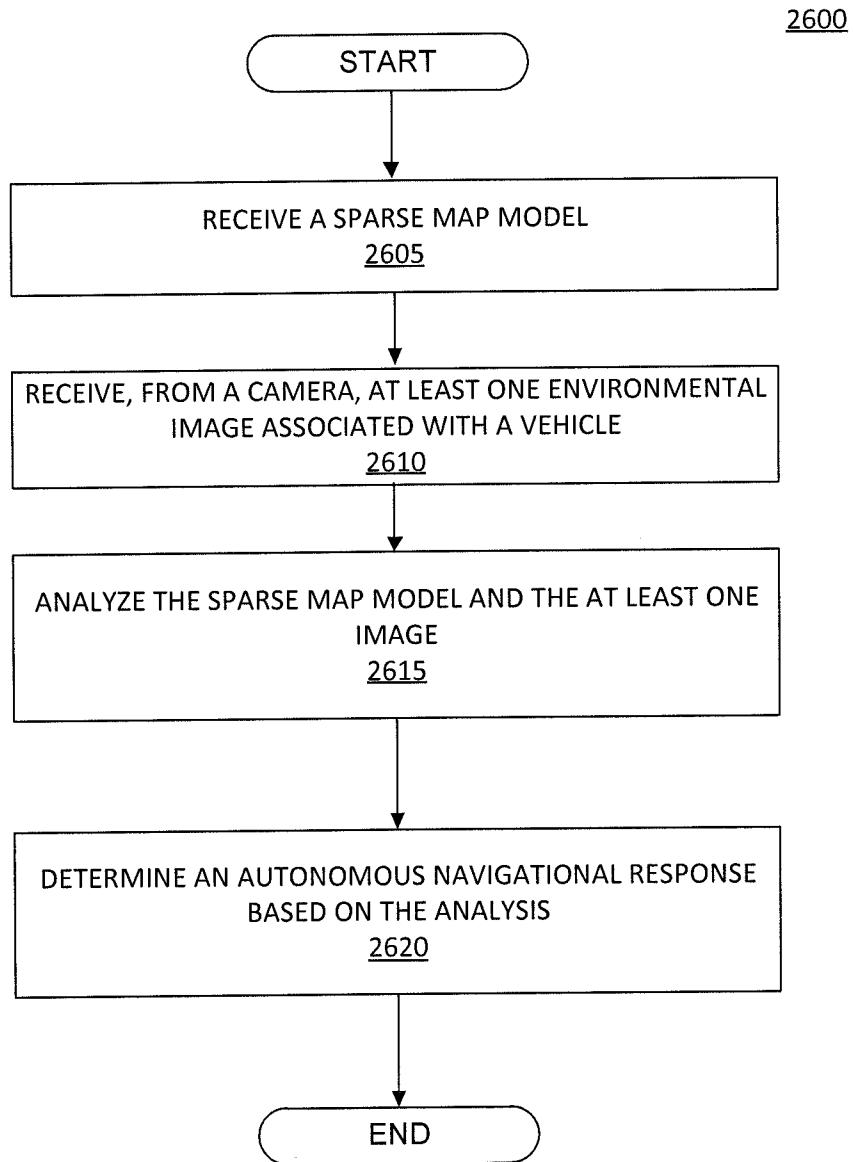


FIG. 26




FIG. 27




FIG. 28A




FIG. 28B


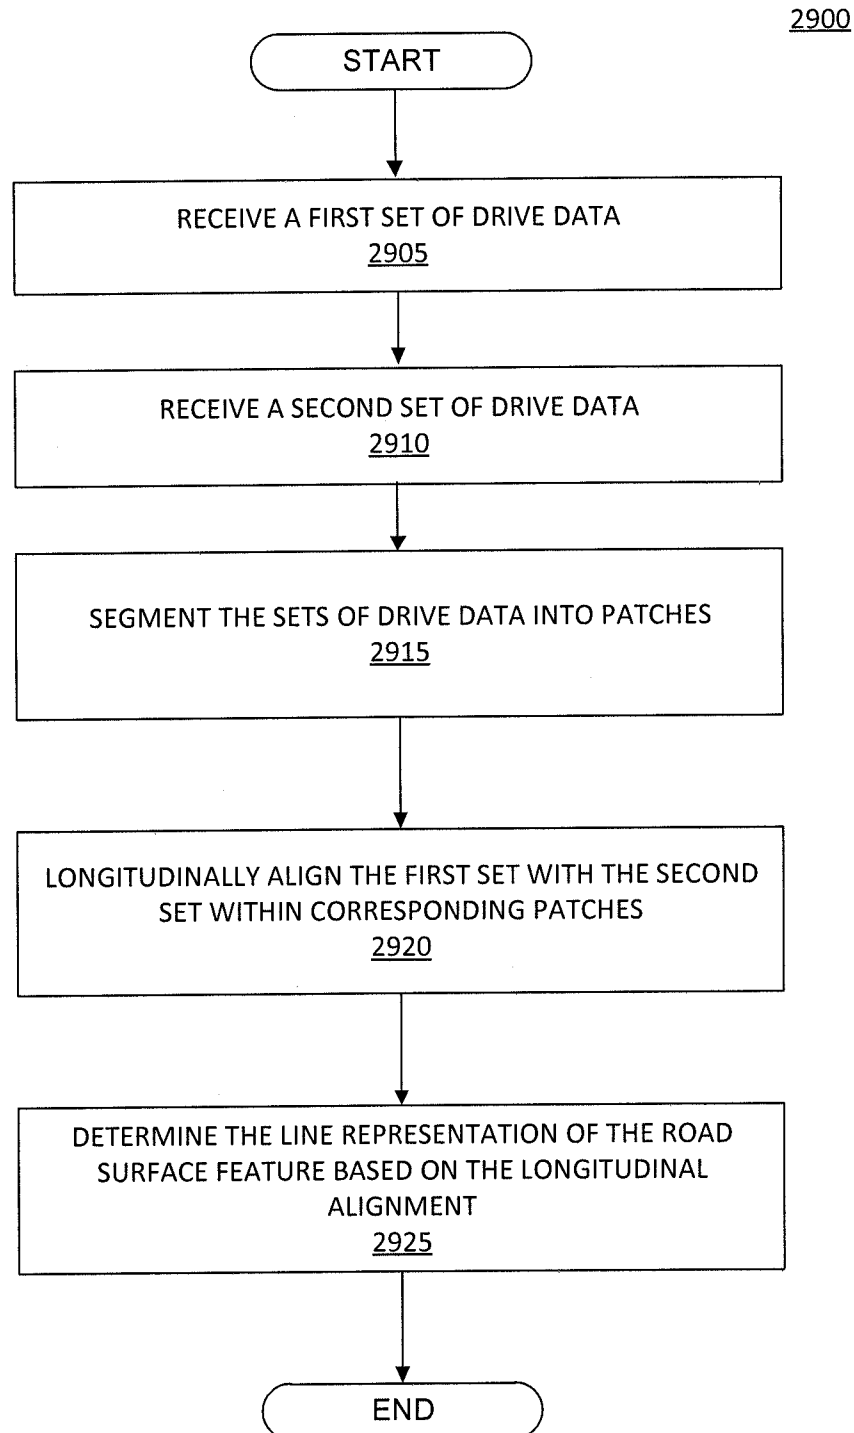


FIG. 28C

47/68

**FIG. 29**


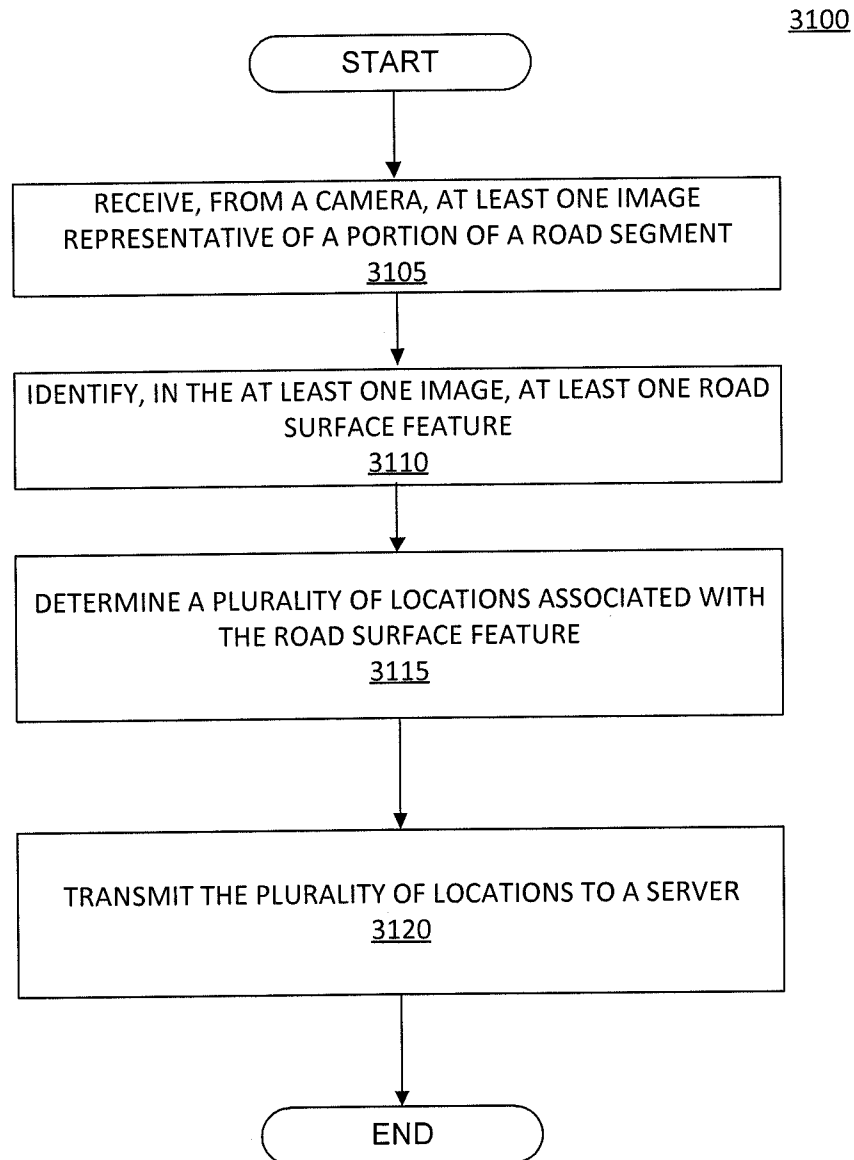


FIG. 30

49/68

**FIG. 31**




FIG. 32




FIG. 33A




FIG. 33B




FIG. 33C




FIG. 34A

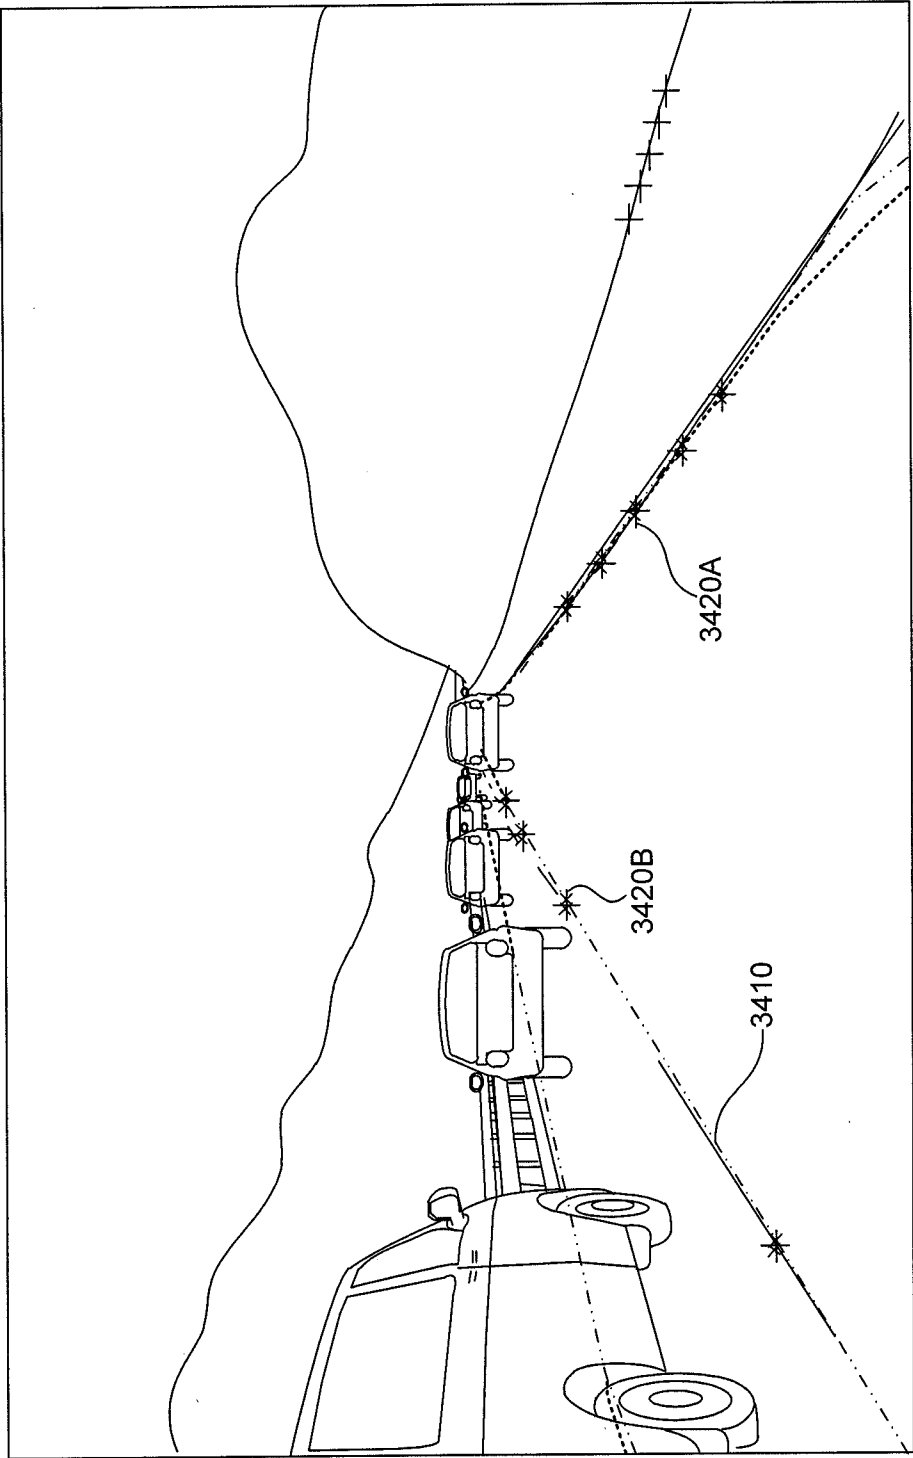


FIG. 34B

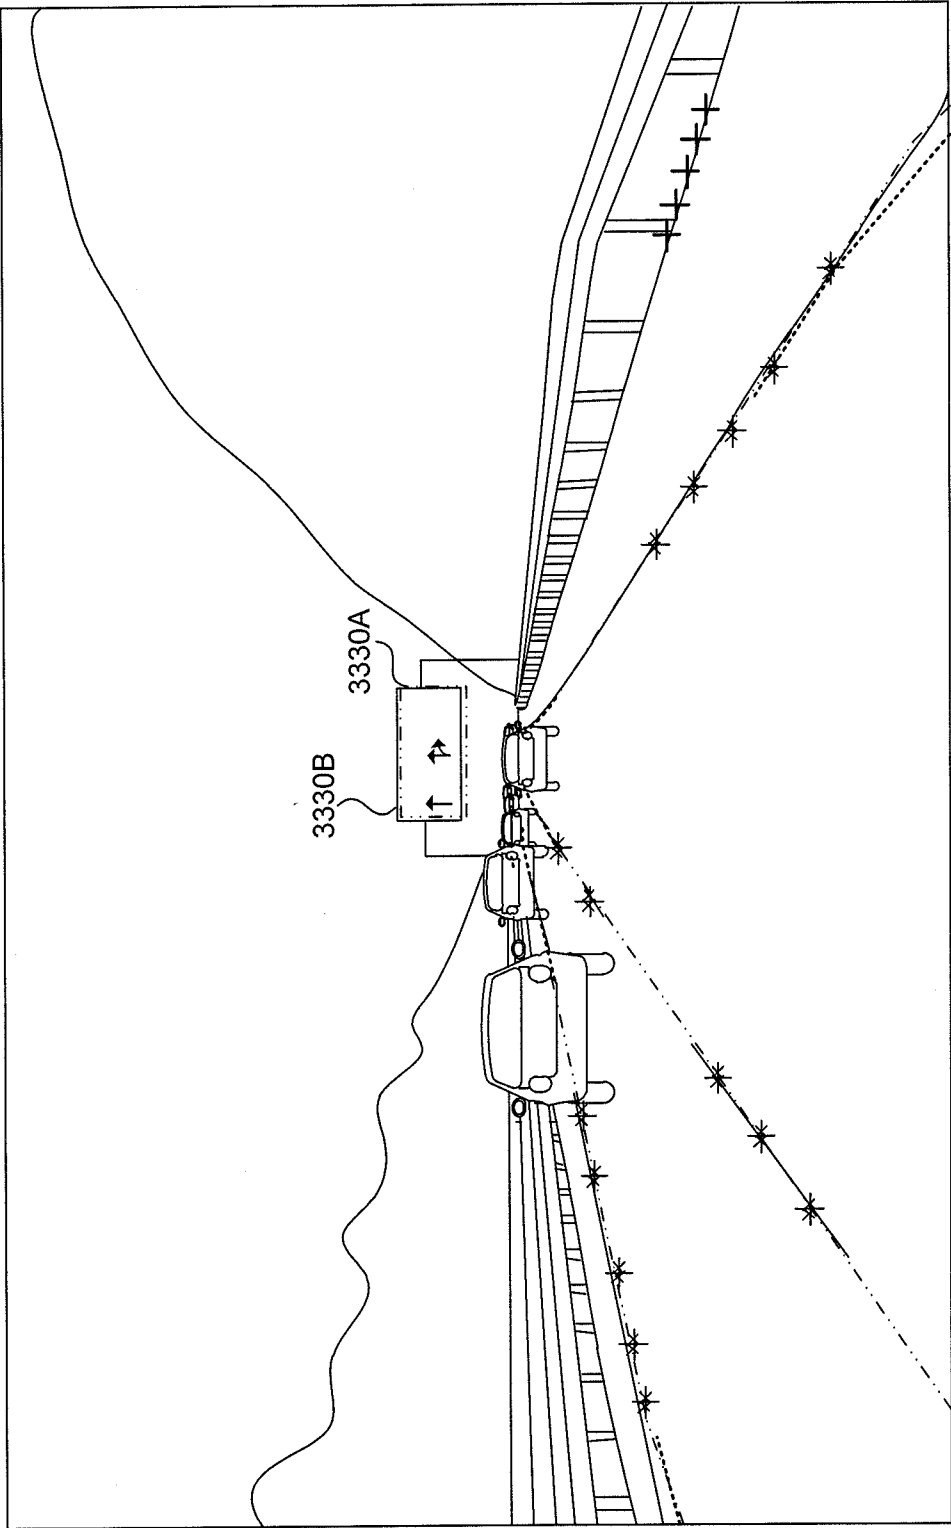
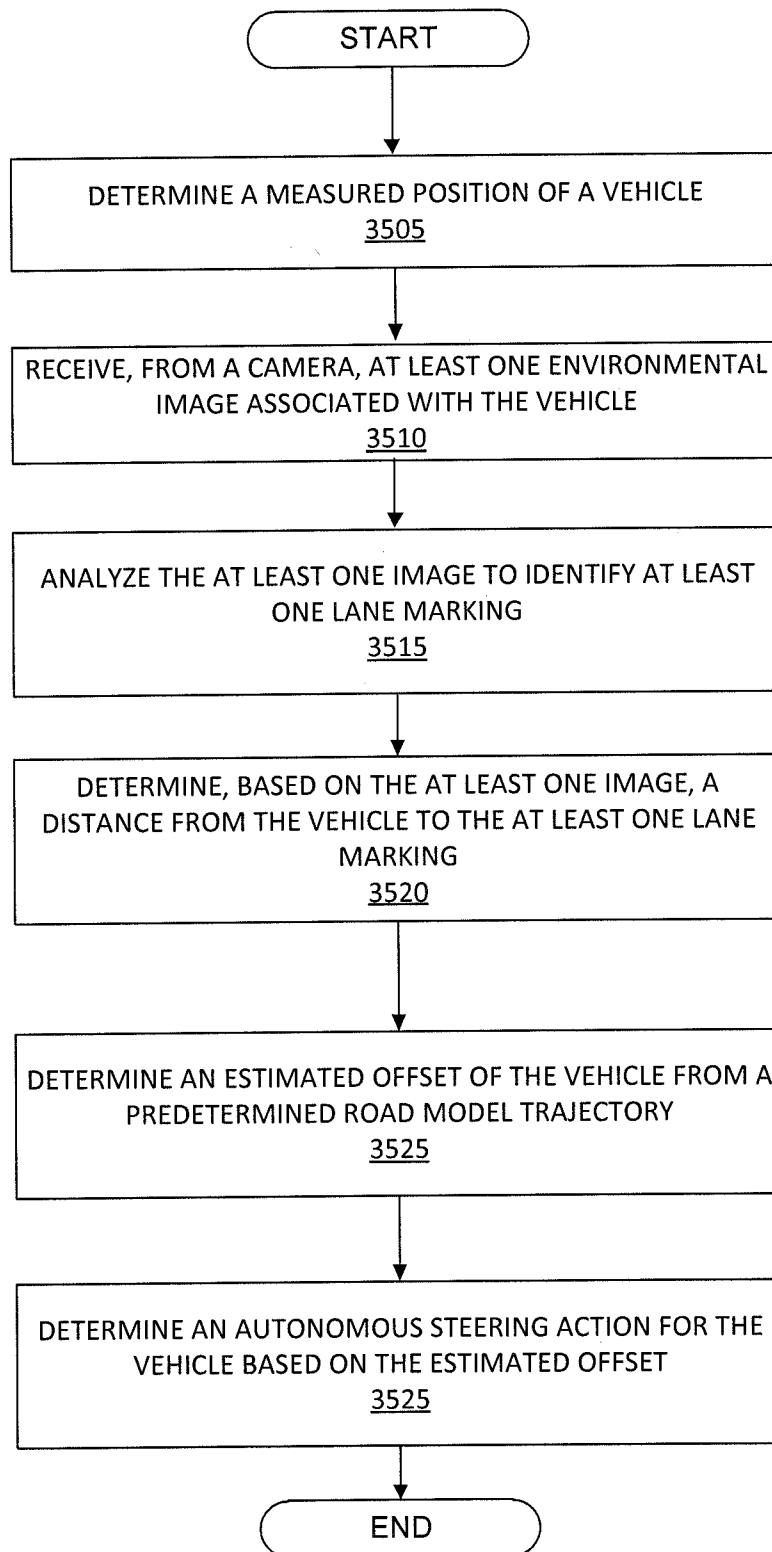


FIG. 34C

57/68

3500**FIG. 35**




FIG. 36




FIG. 37A




FIG. 37B




FIG. 37C




FIG. 38A




FIG. 38B

62/68




FIG. 39A




FIG. 39B




FIG. 39C




FIG. 40A




FIG. 40B




FIG. 40C


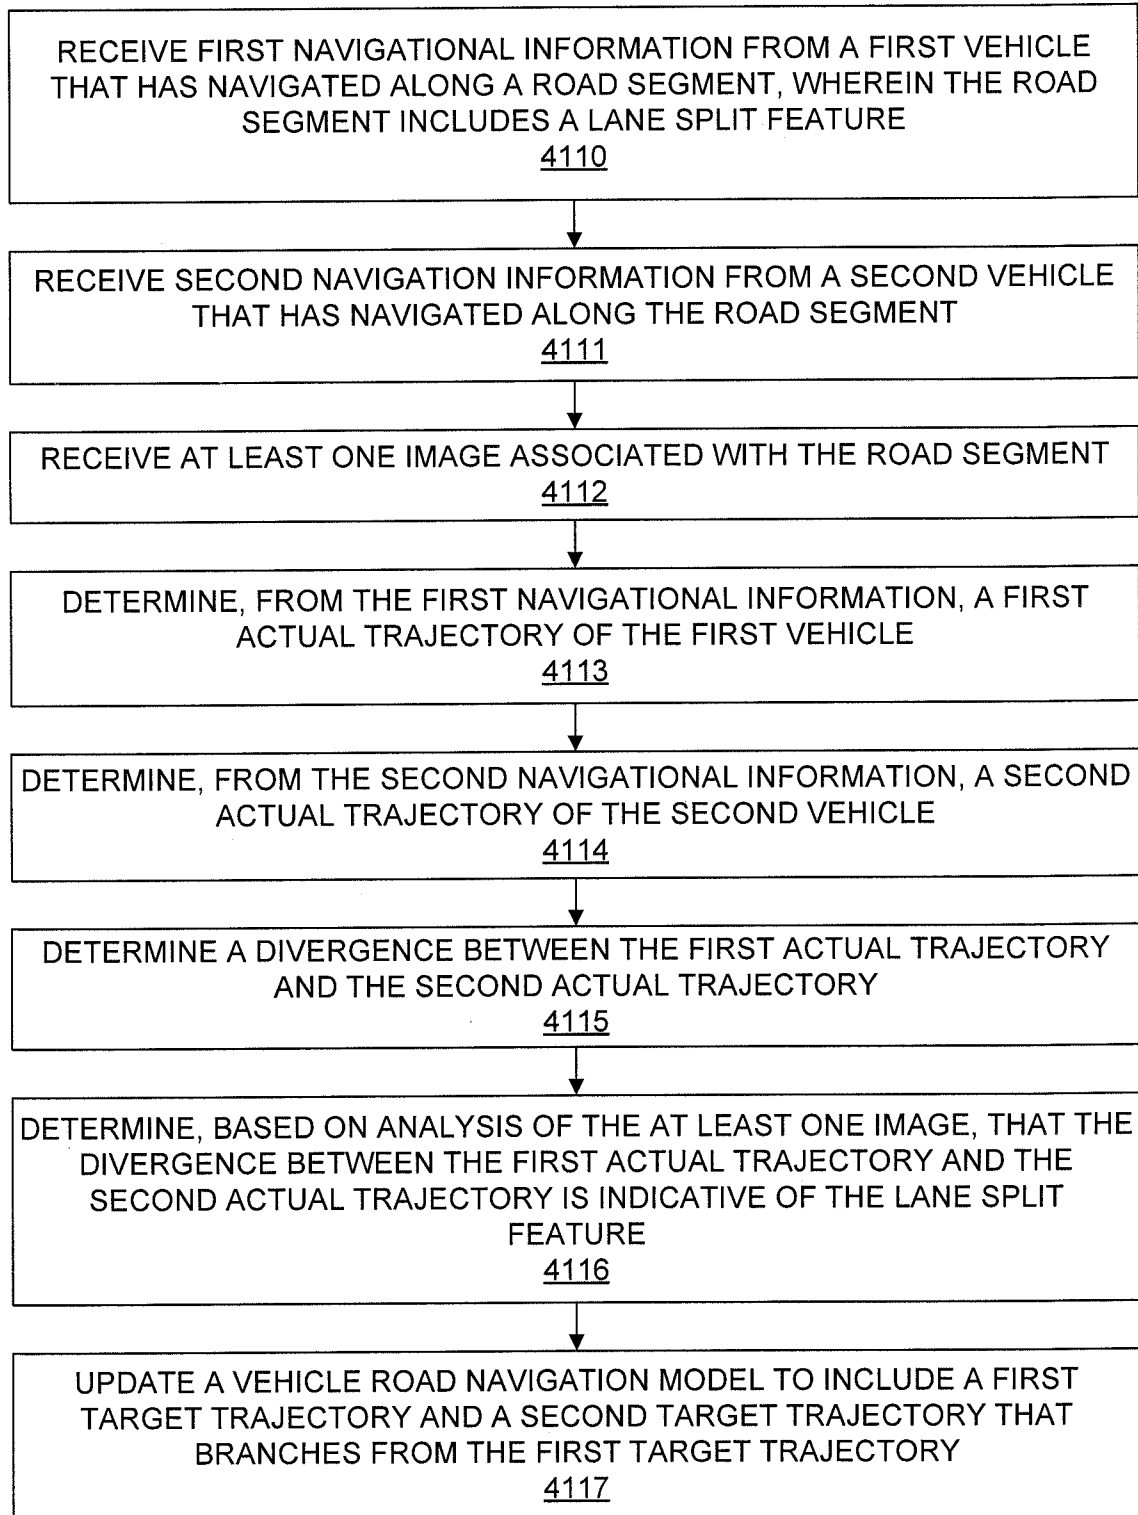
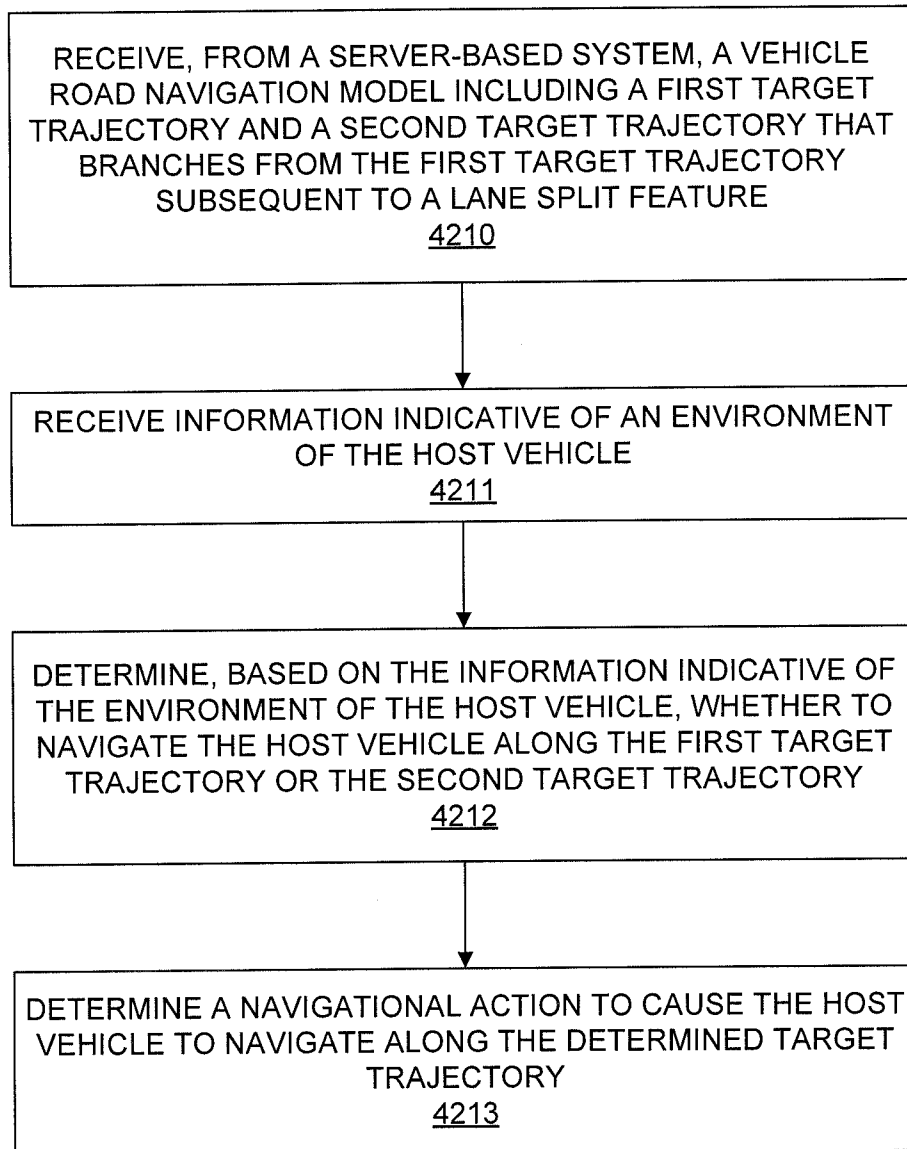


FIG. 40D

67/68

4100**FIG. 41**

4200**FIG. 42**