



US 20190034740A1

(19) **United States**

(12) **Patent Application Publication**  
**KWANT et al.**

(10) **Pub. No.: US 2019/0034740 A1**

(43) **Pub. Date: Jan. 31, 2019**

(54) **METHOD, APPARATUS, AND SYSTEM FOR VANISHING POINT/HORIZON ESTIMATION USING LANE MODELS**

(52) **U.S. Cl.**  
CPC ..... **G06K 9/00798** (2013.01); **G06K 9/4671** (2013.01)

(71) Applicant: **HERE GLOBAL B.V.**, Eindhoven (NL)

(57) **ABSTRACT**

(72) Inventors: **Richard KWANT**, Oakland, CA (US);  
**Anish MITTAL**, Berkeley, CA (US);  
**Nicholas POJMAN**, San Francisco, CA (US);  
**Yangyang CHEN**, Oakland, CA (US)

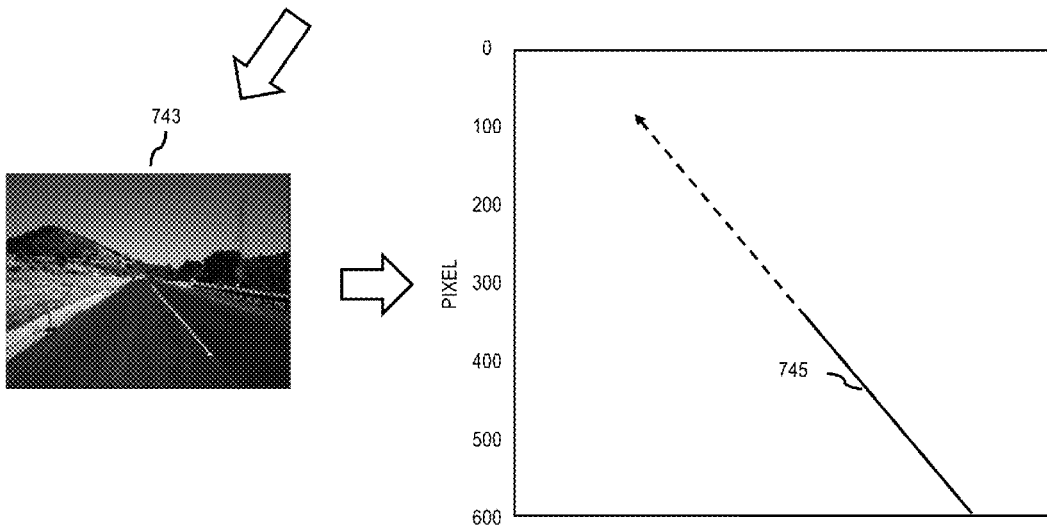
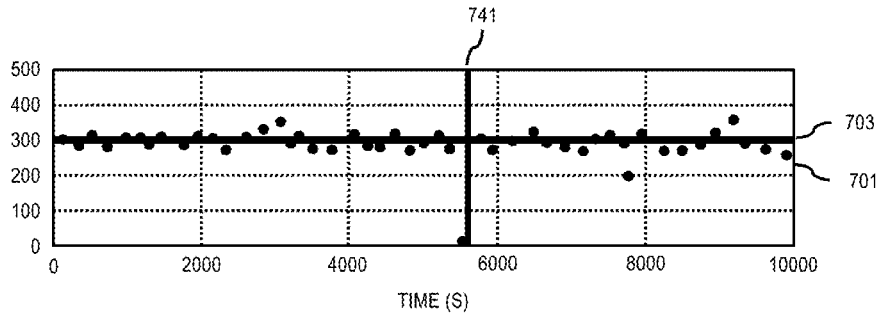
An approach is provided for estimating a vanishing point or horizon in an image depicting one or more lanes of a roadway. The approach involves processing the image to construct one or more lane models of the one or more road lanes depicted in the image. The approach also involves extending the one or more road lanes through the image using the one or more lane models. The approach further involves determining a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line. The approach further involves designating the horizontal line as the horizon of the image.

(21) Appl. No.: **15/661,915**

(22) Filed: **Jul. 27, 2017**

**Publication Classification**

(51) **Int. Cl.**  
**G06K 9/00** (2006.01)  
**G06K 9/46** (2006.01)



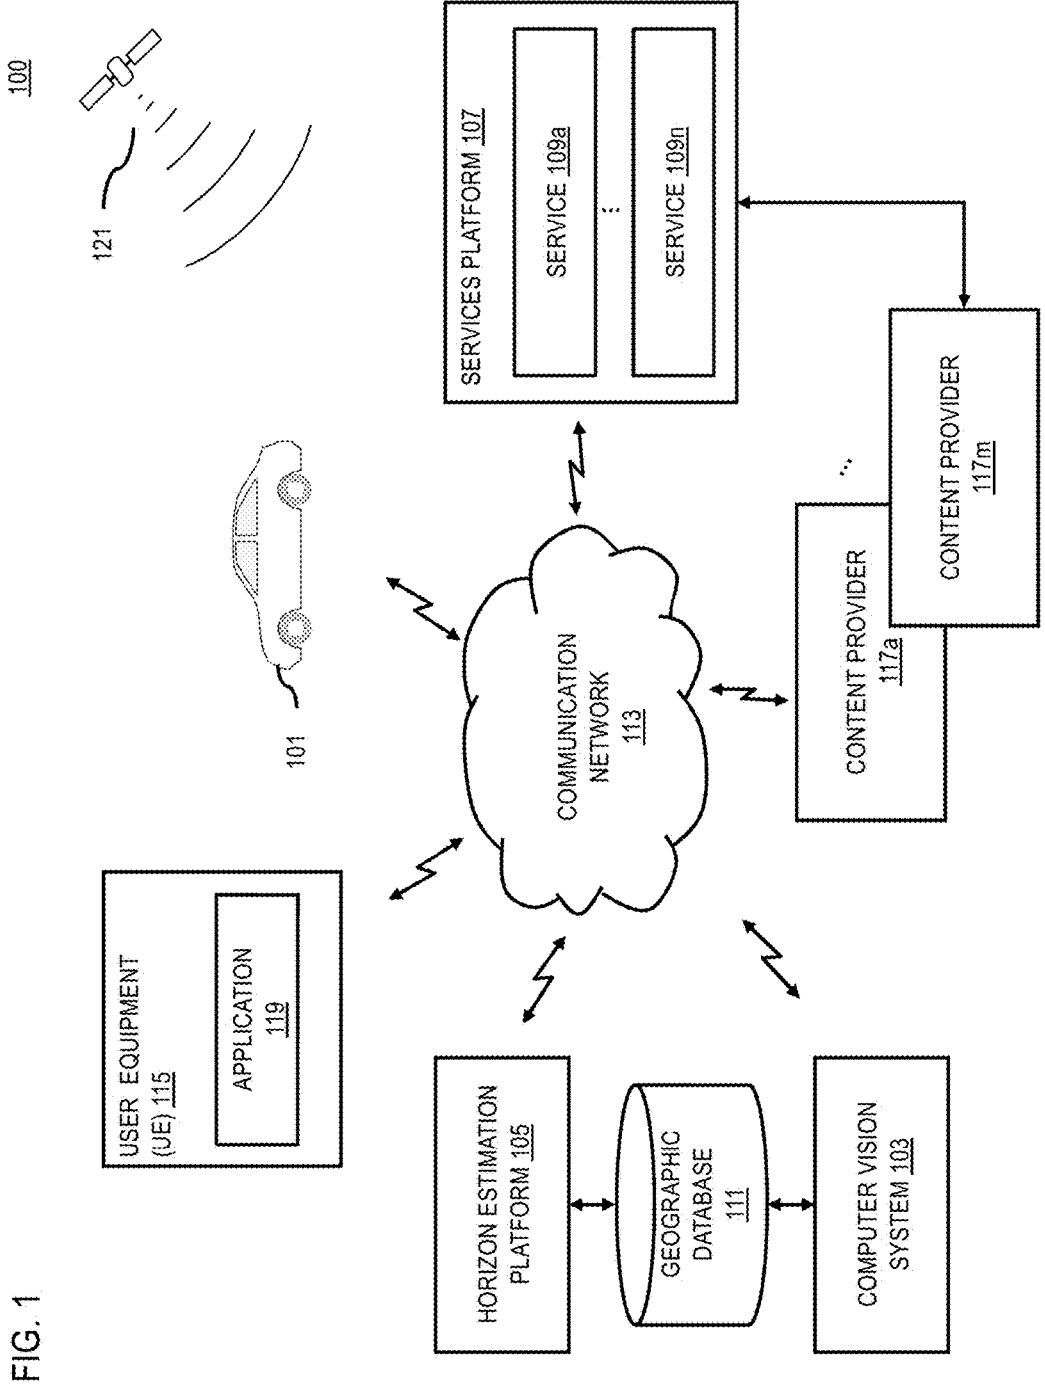


FIG. 1

FIG. 2

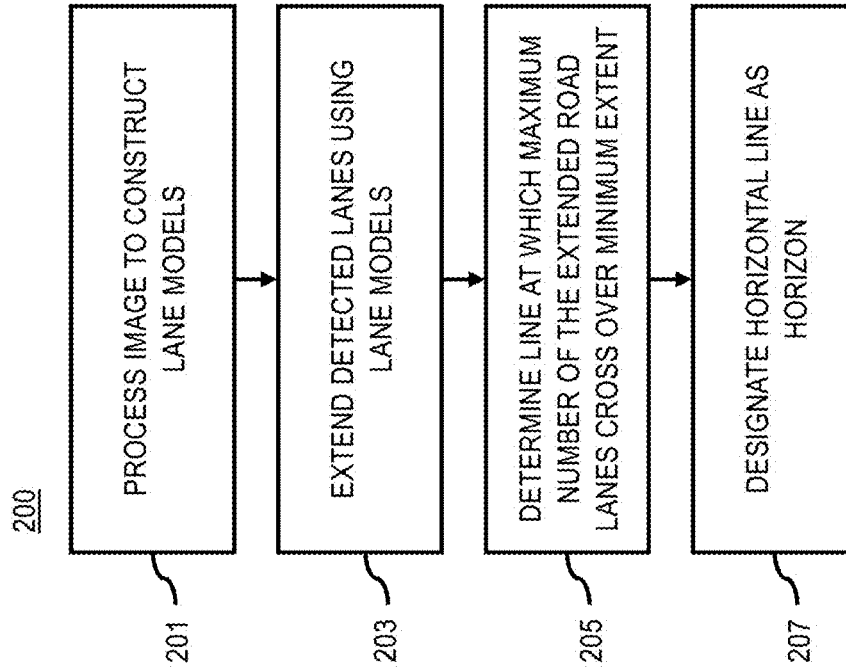


FIG. 3

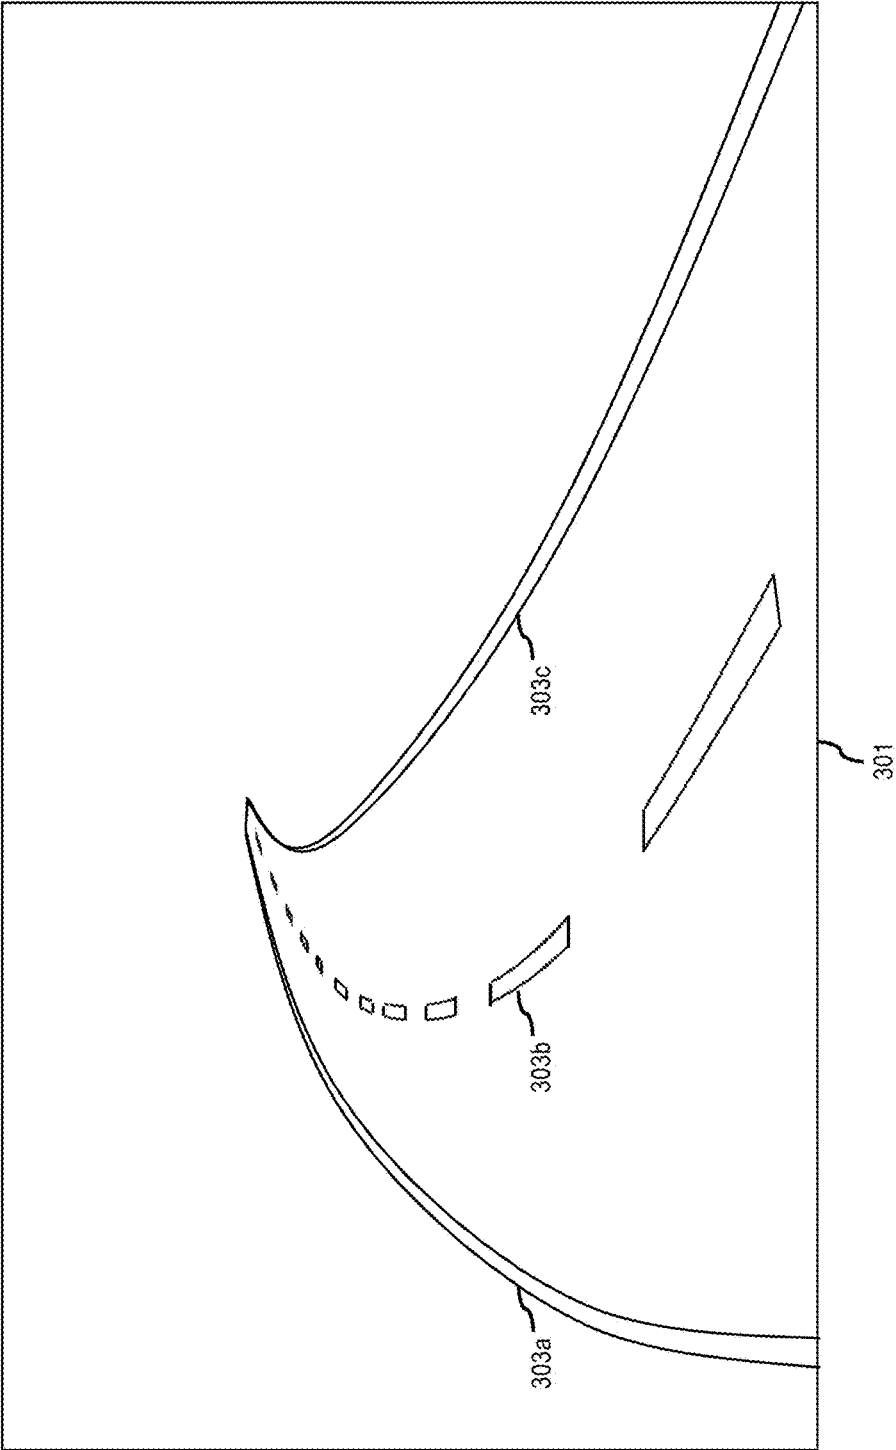


FIG. 4A

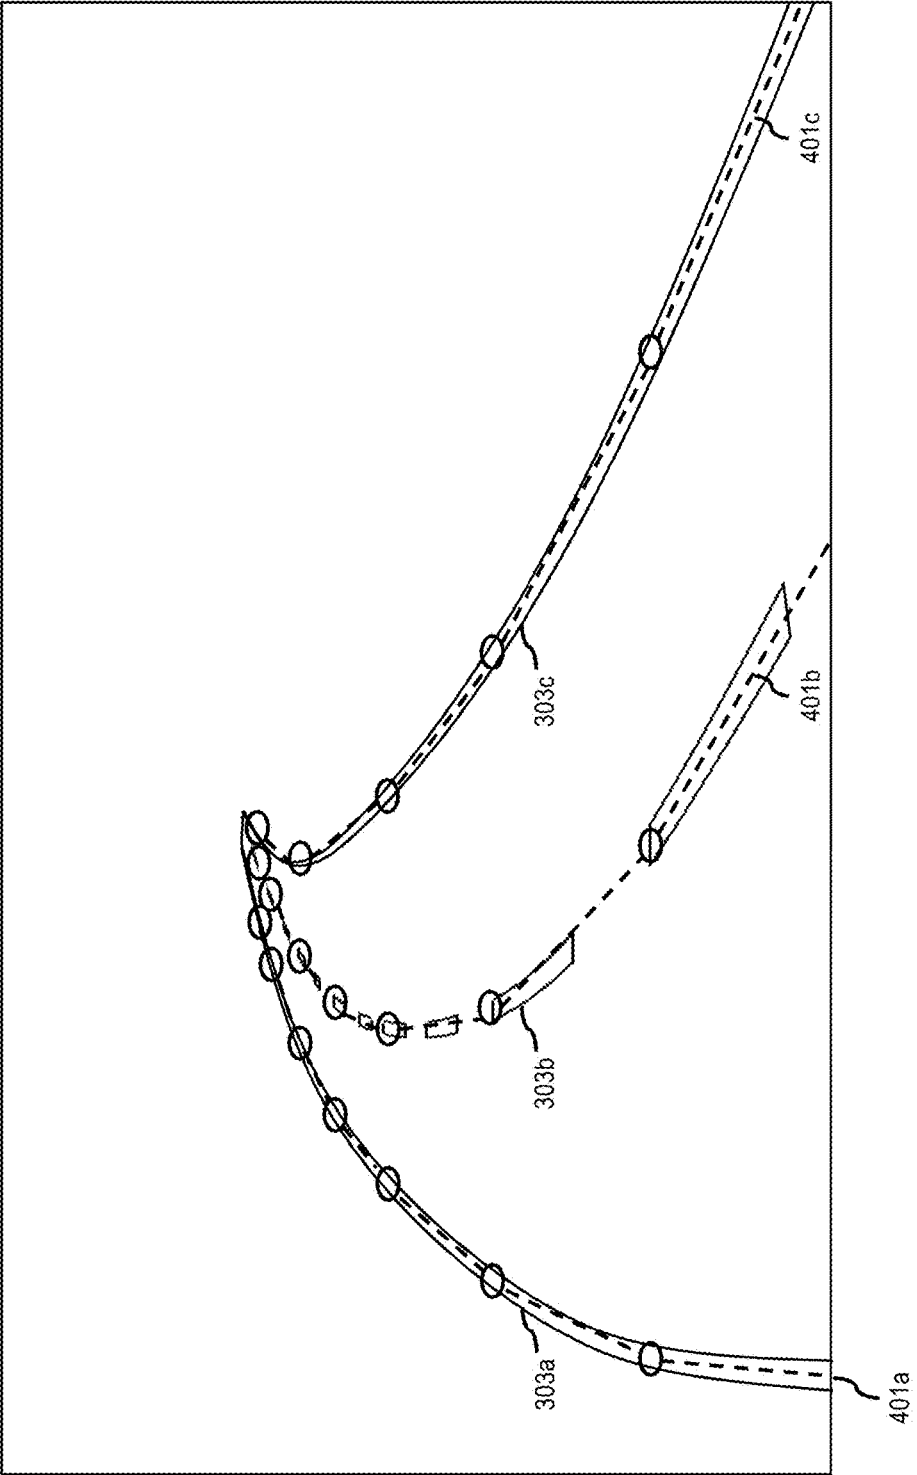


FIG. 4B

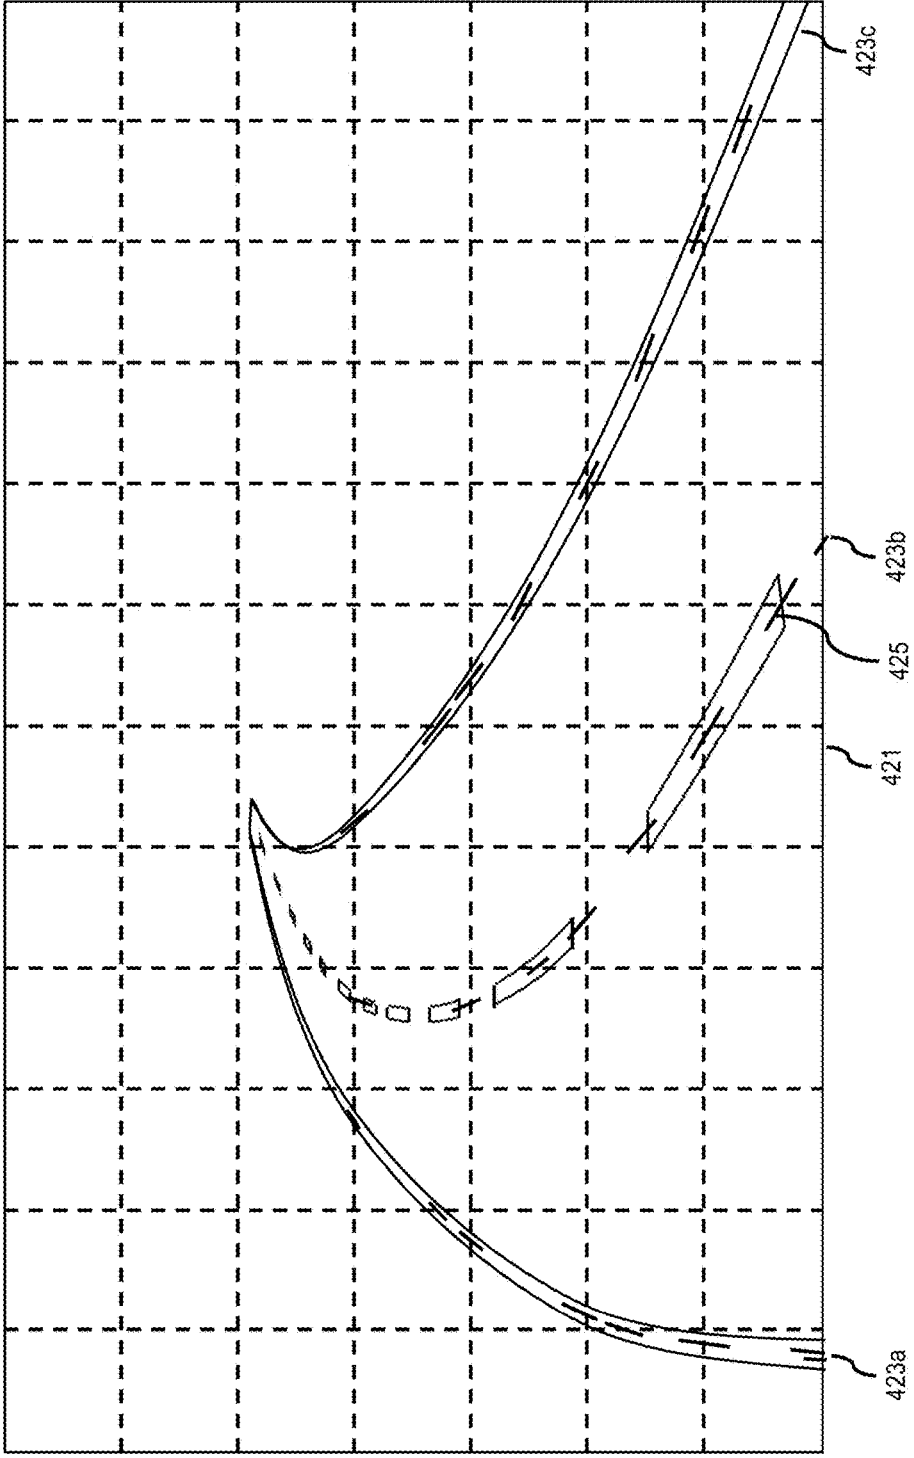


FIG. 5A

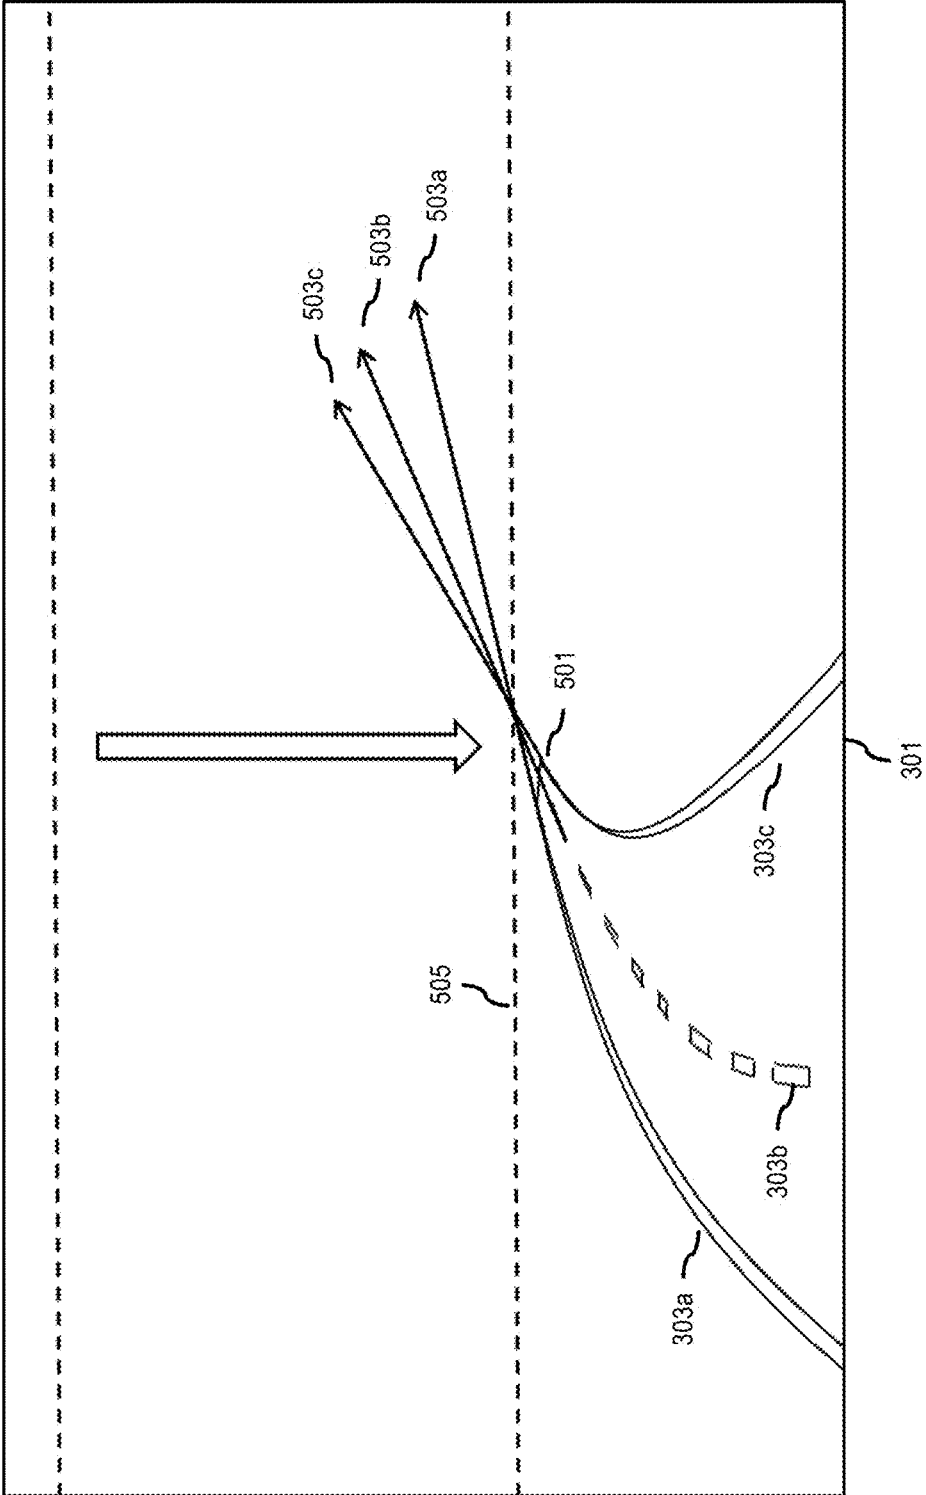
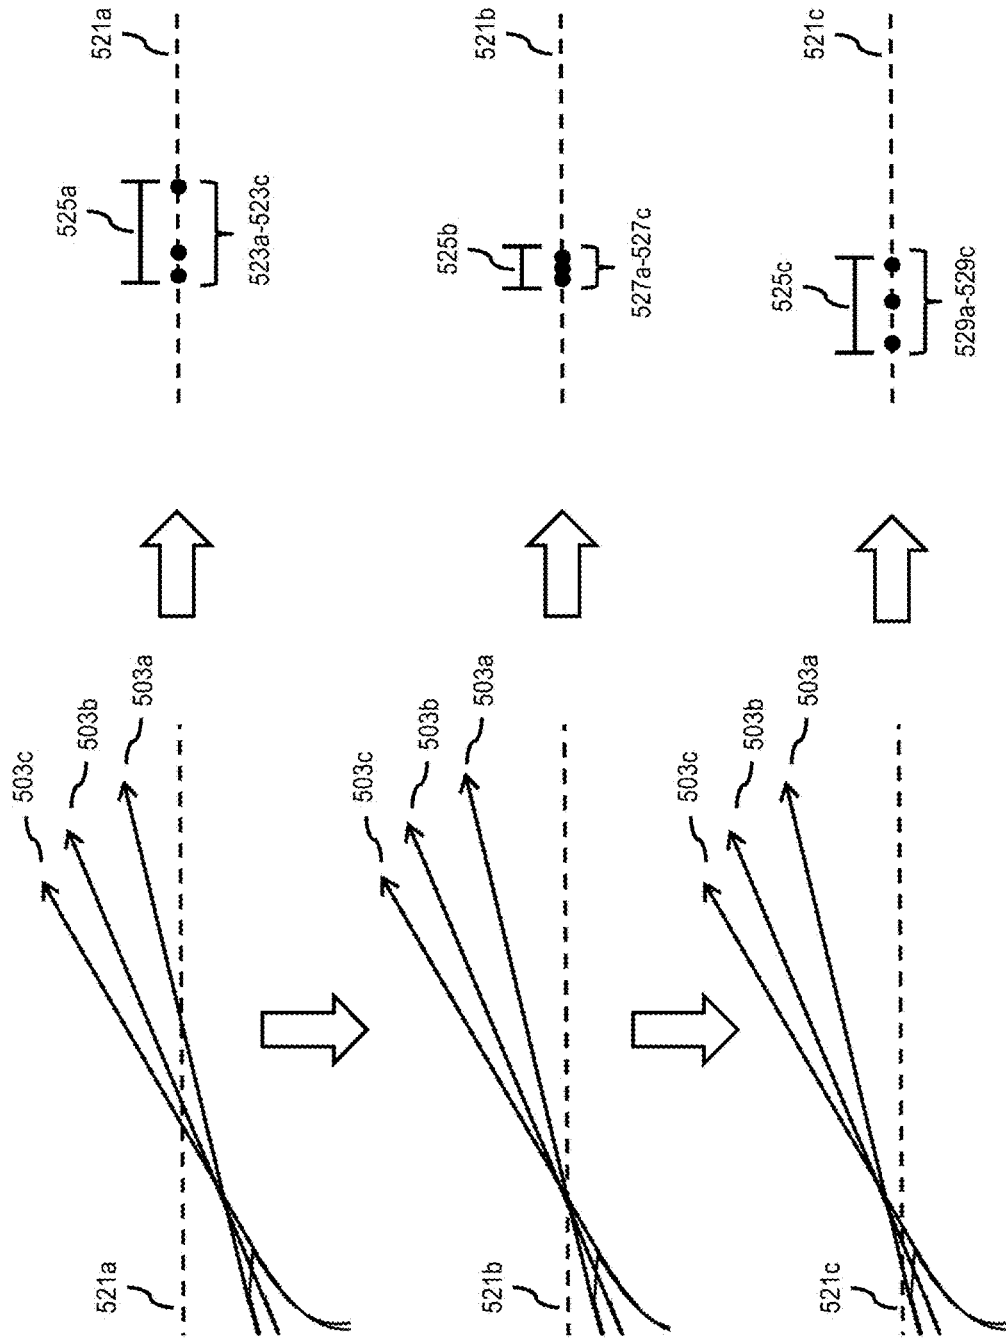


FIG. 5B



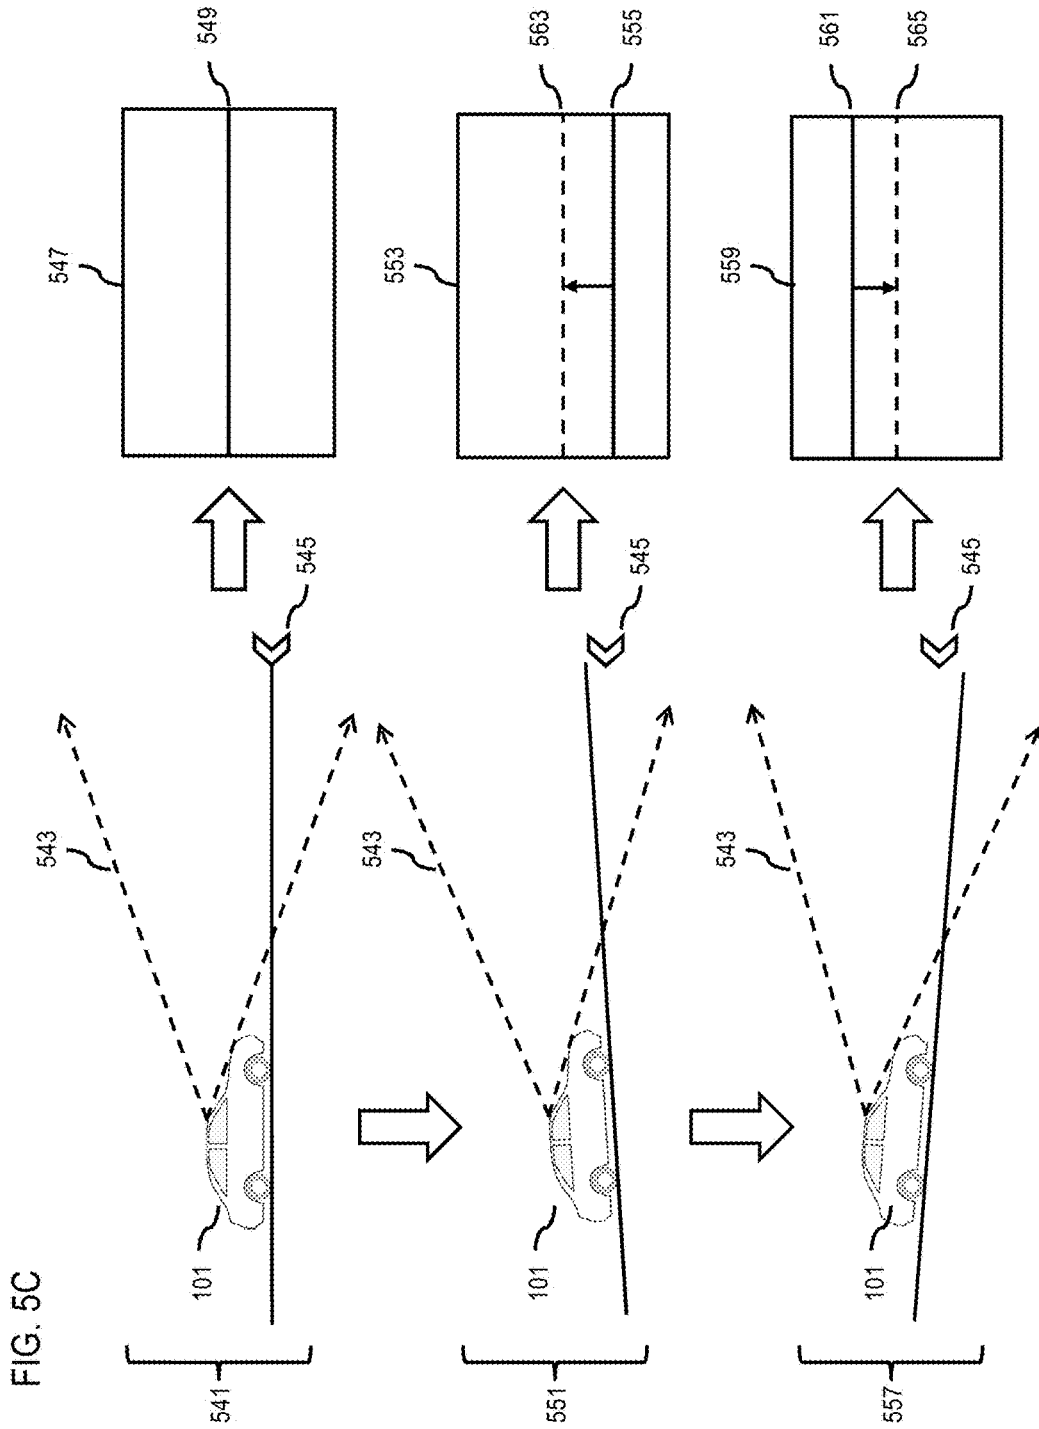
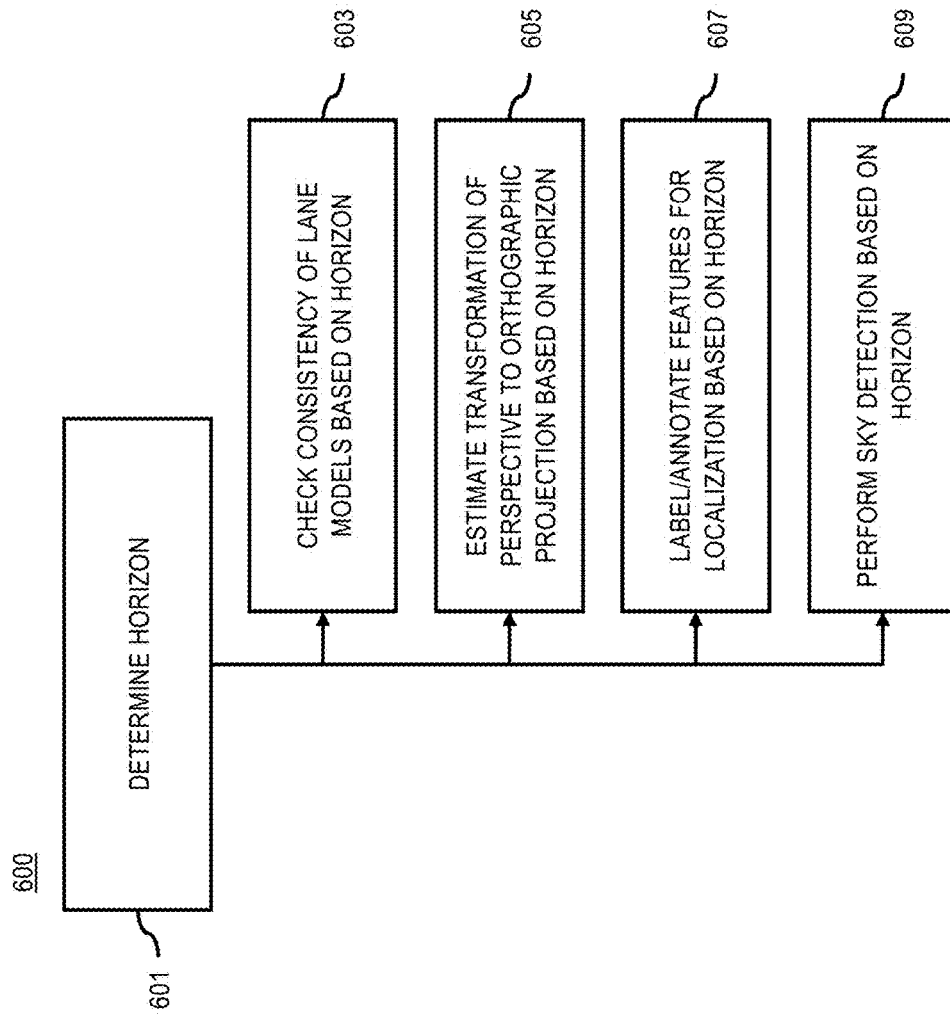
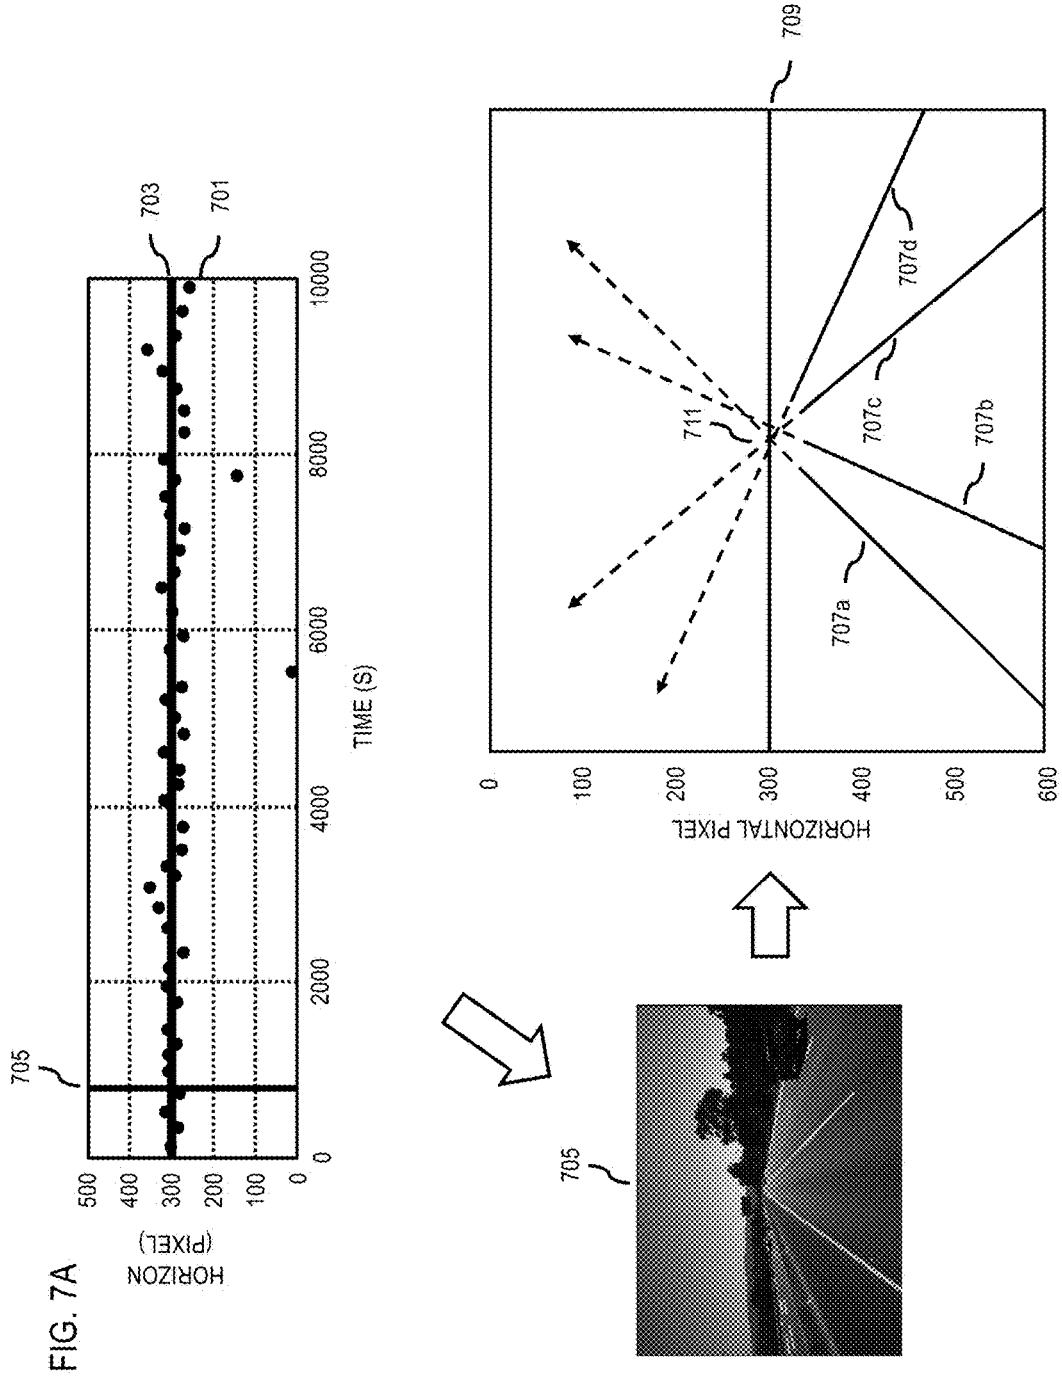


FIG. 6





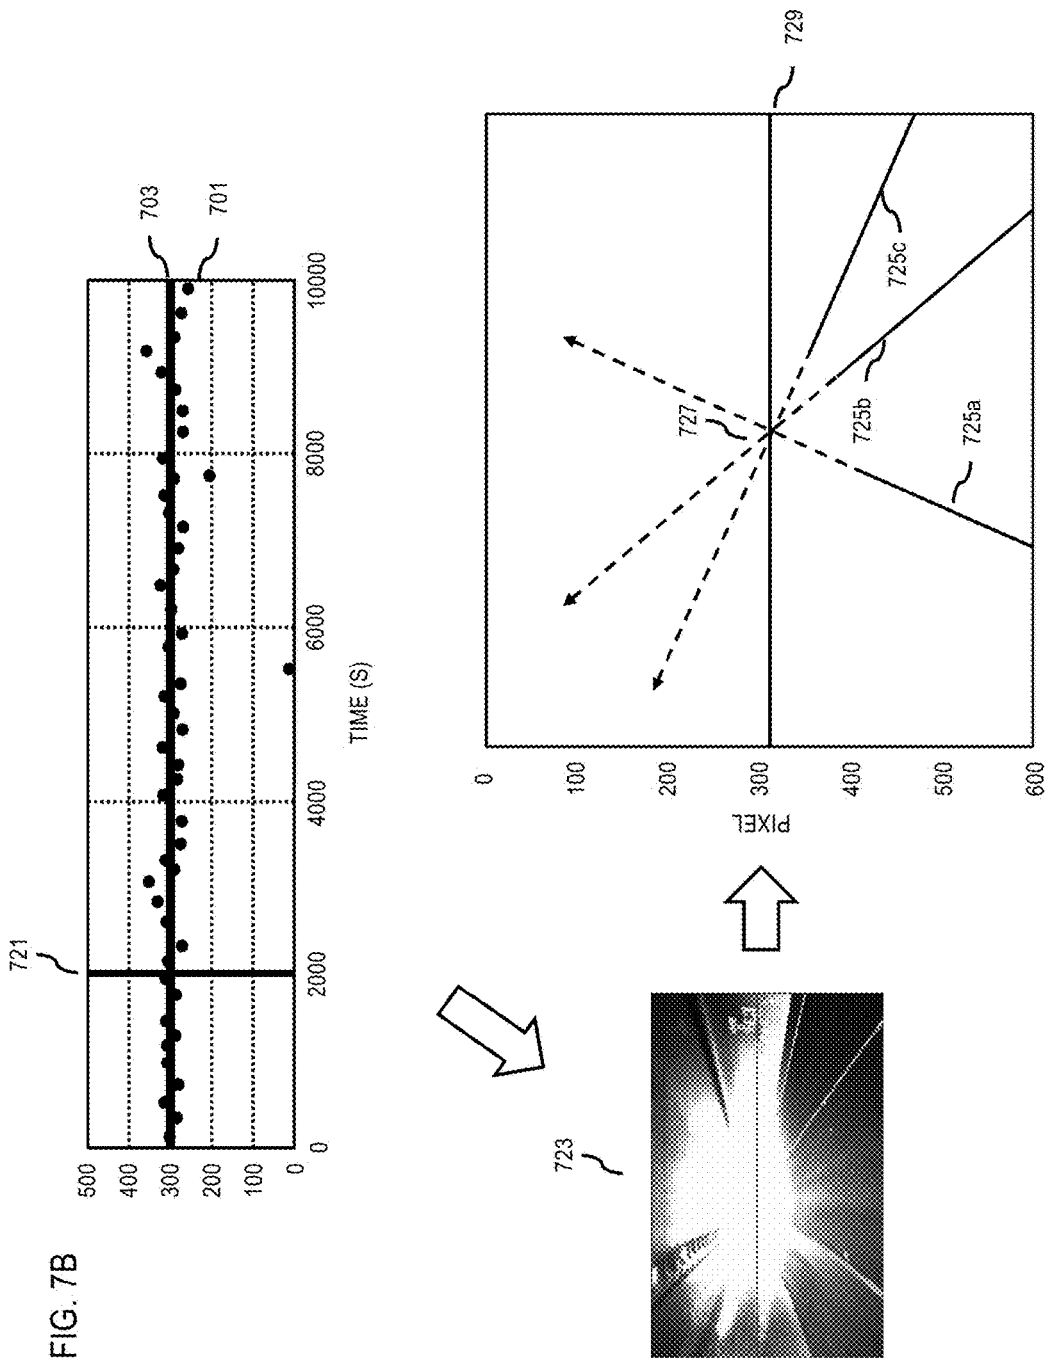


FIG. 7B

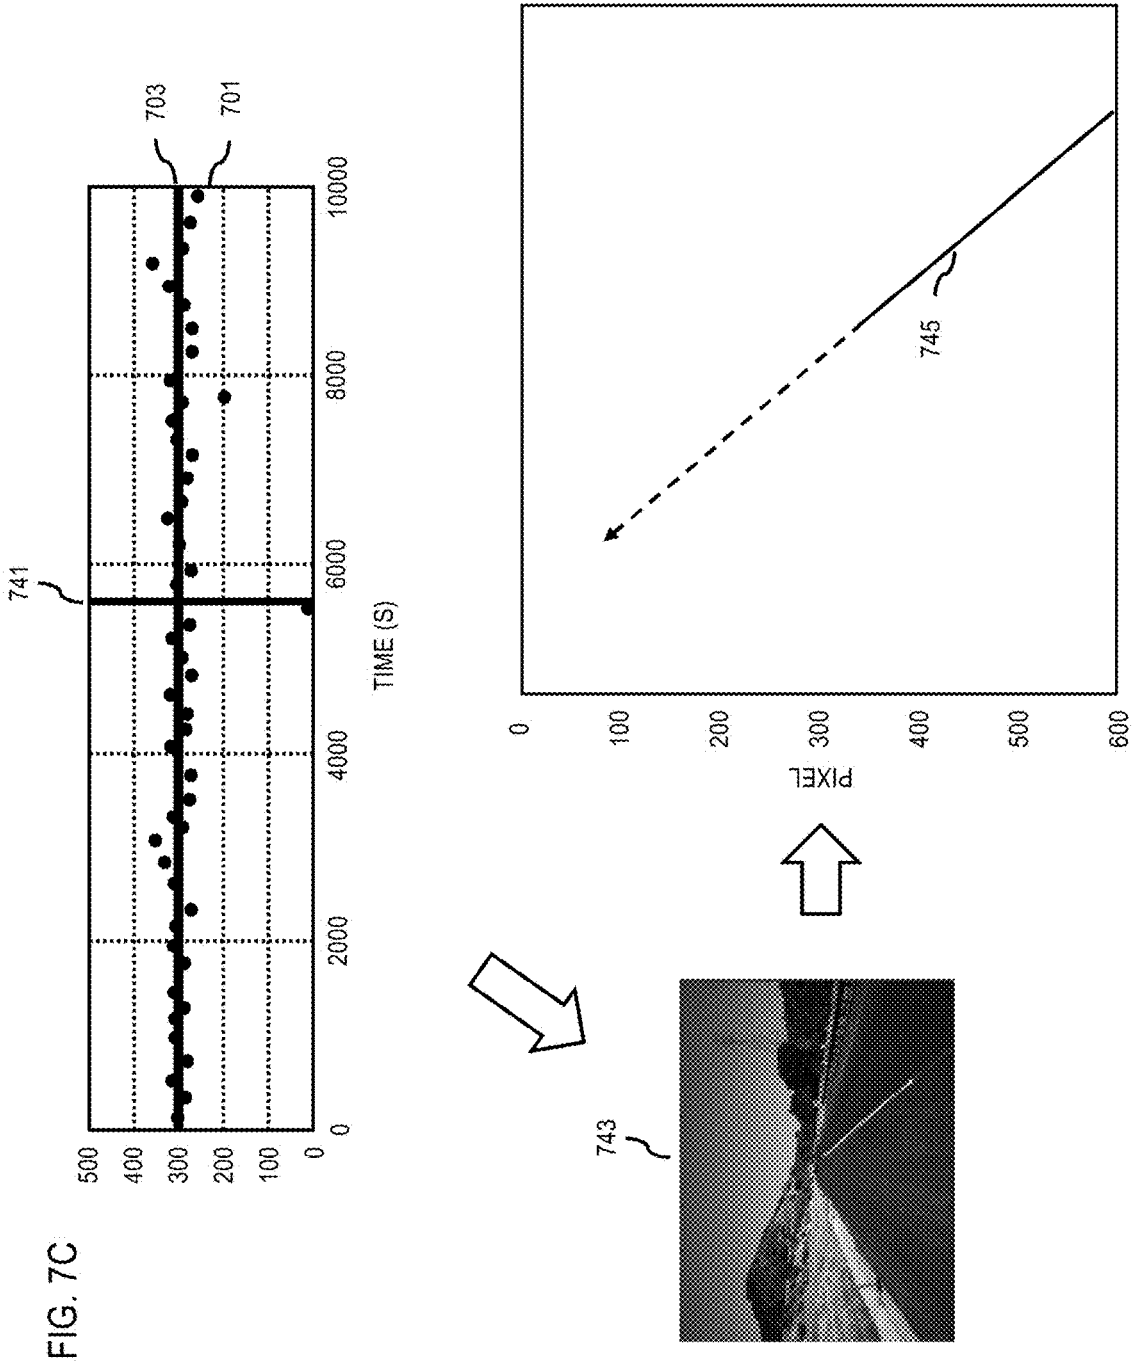


FIG. 7D

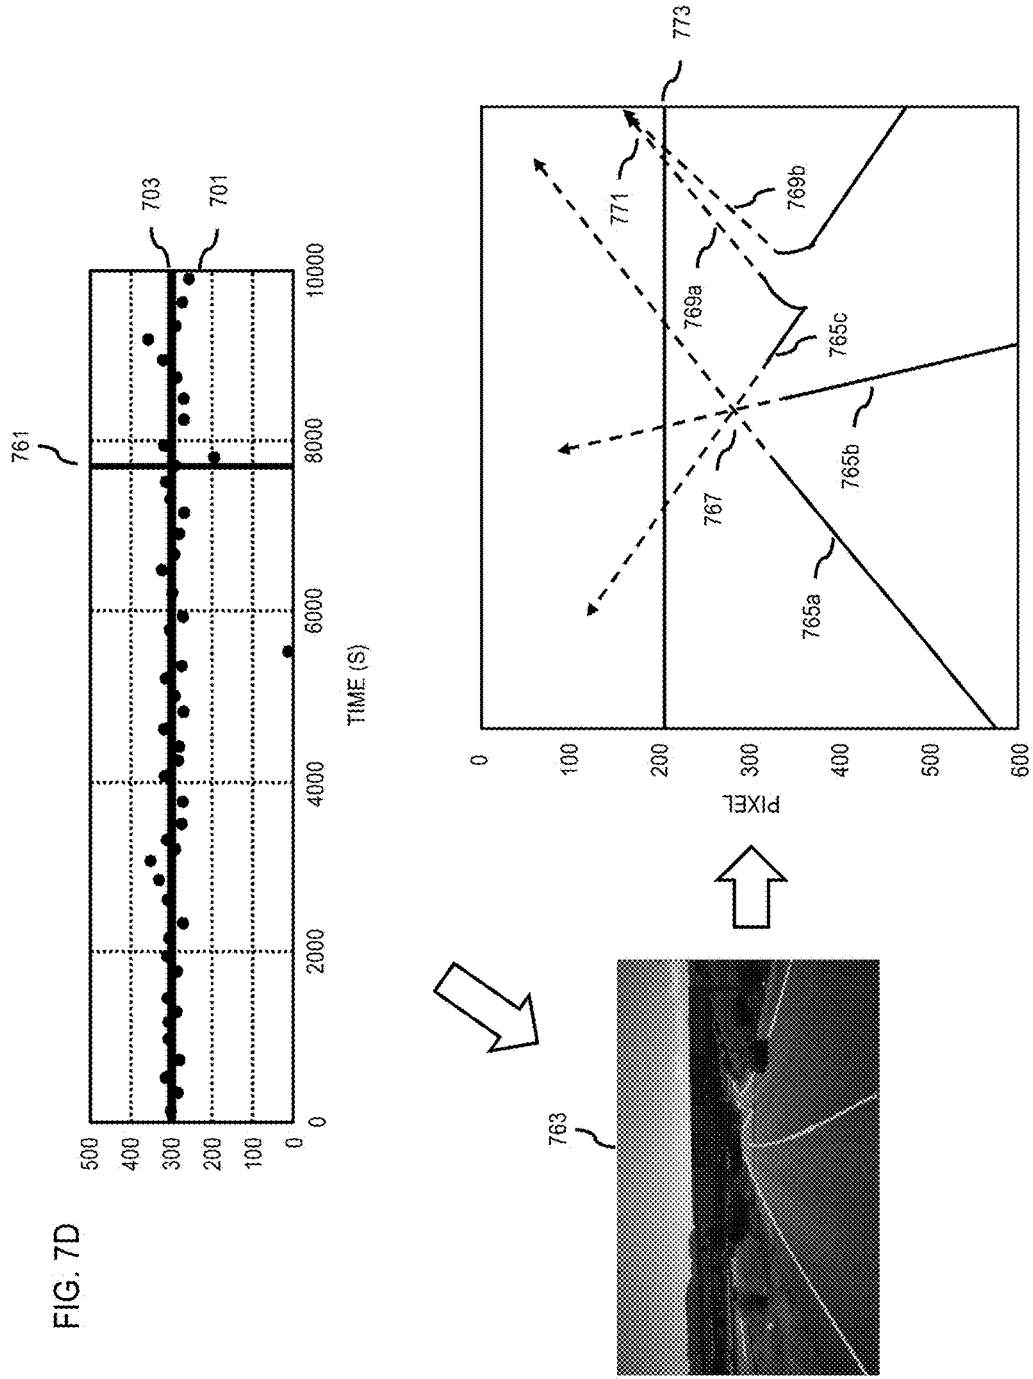
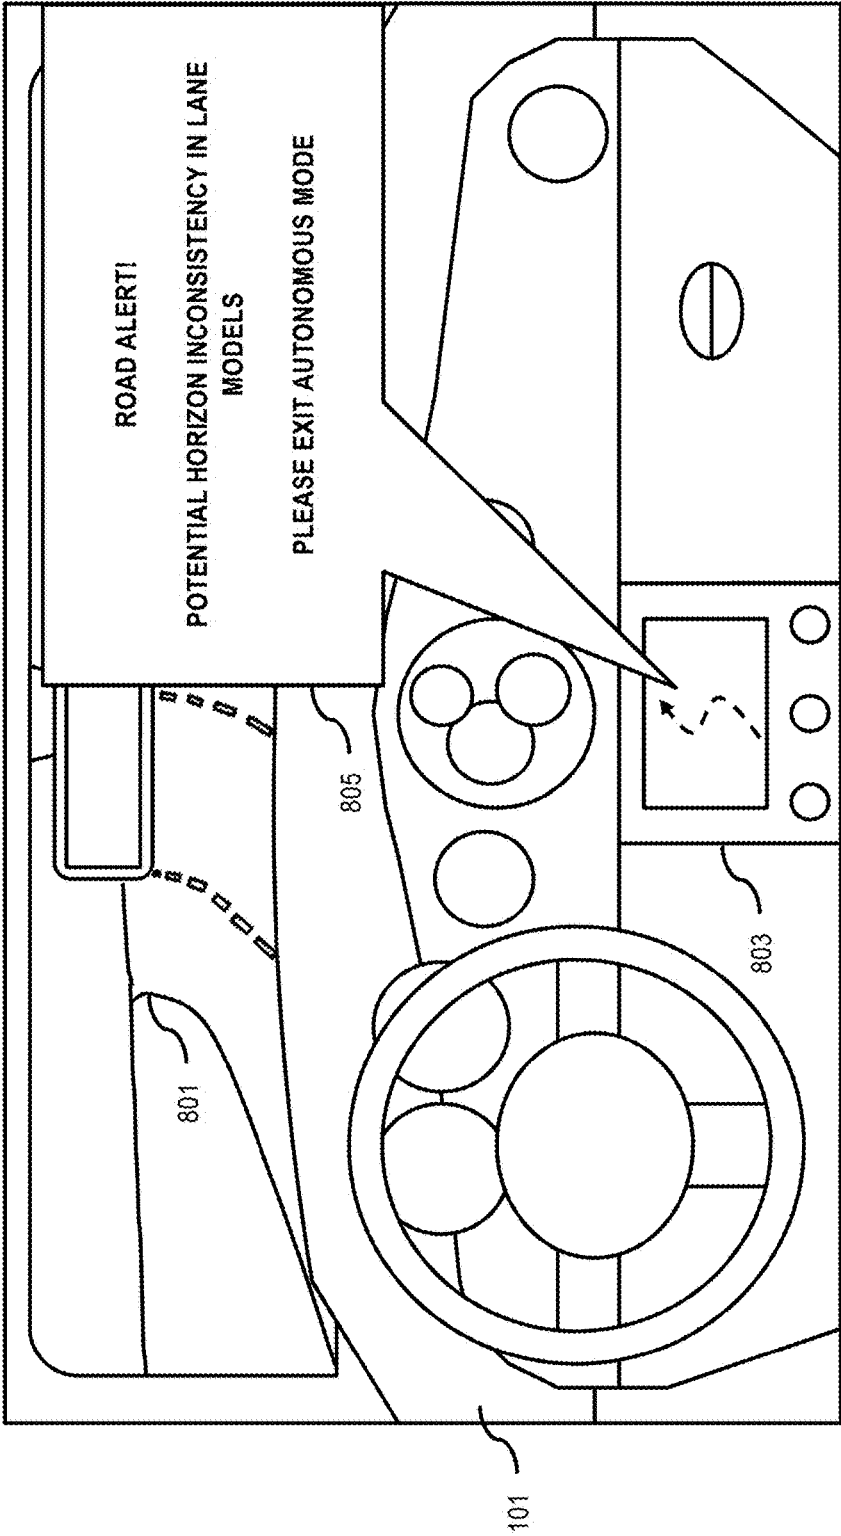


FIG. 8A



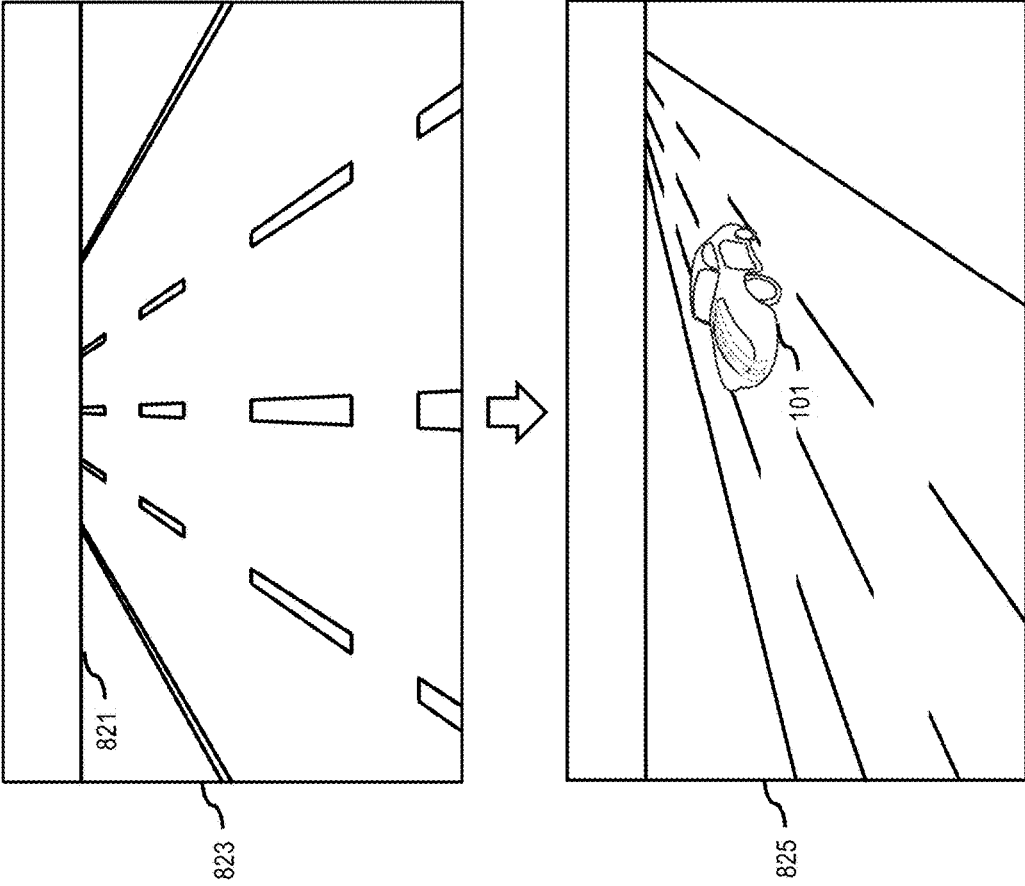


FIG. 8B

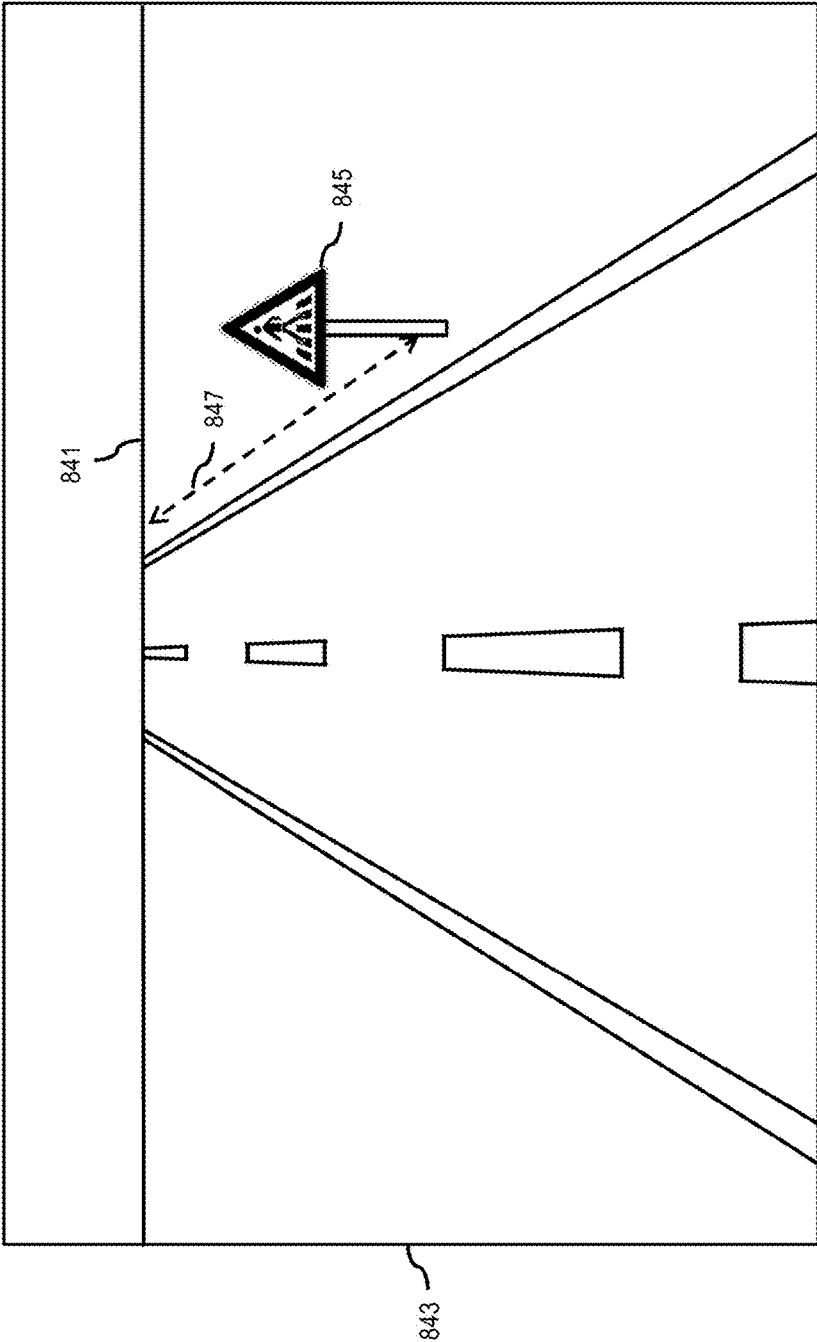


FIG. 8C

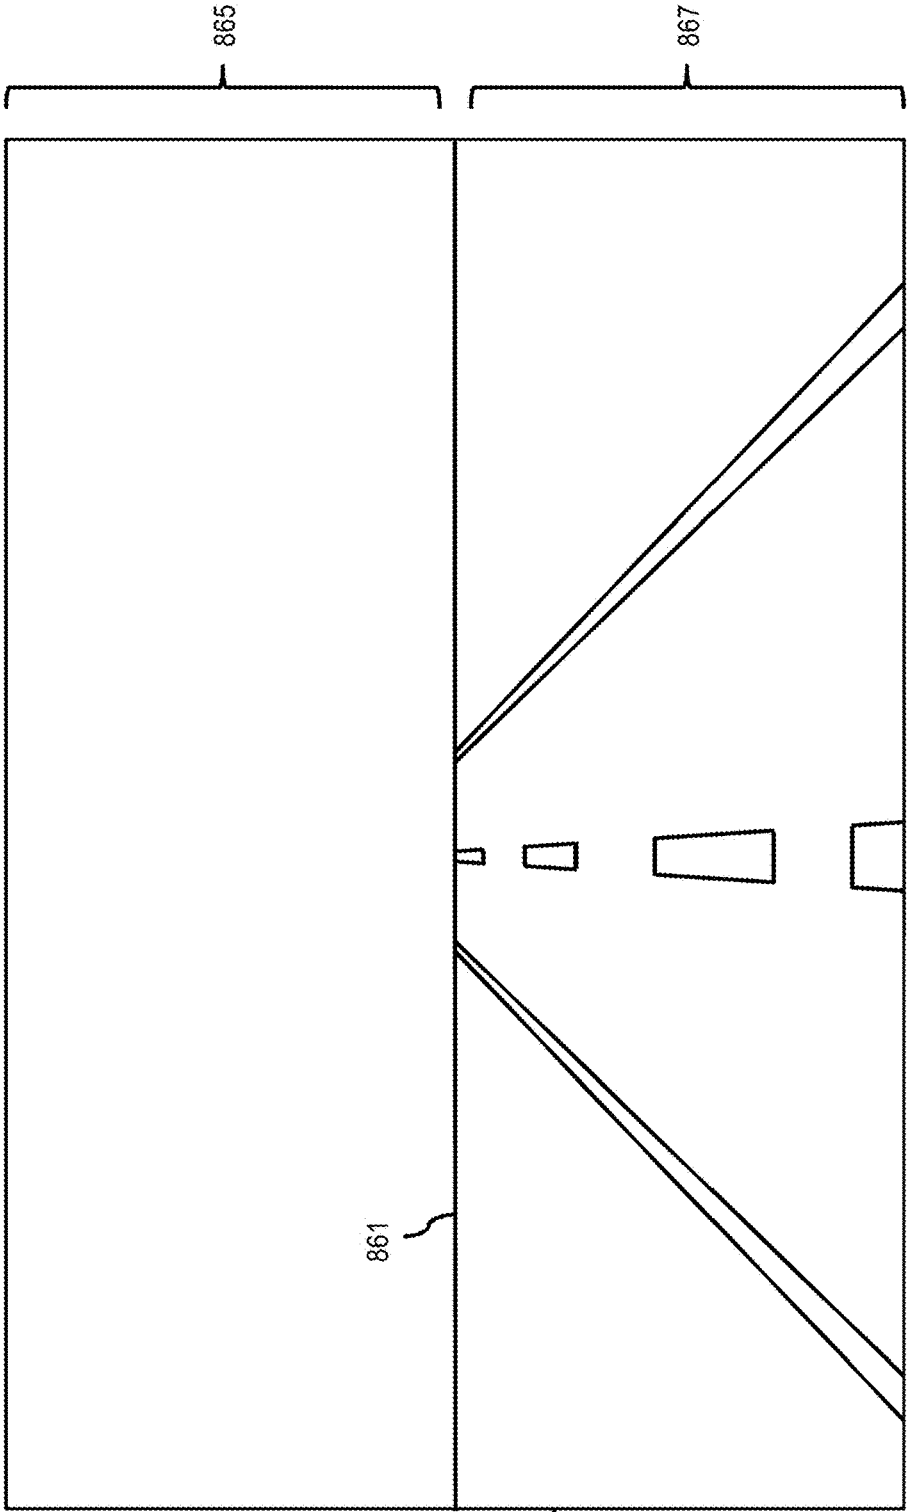


FIG. 8D

FIG. 9

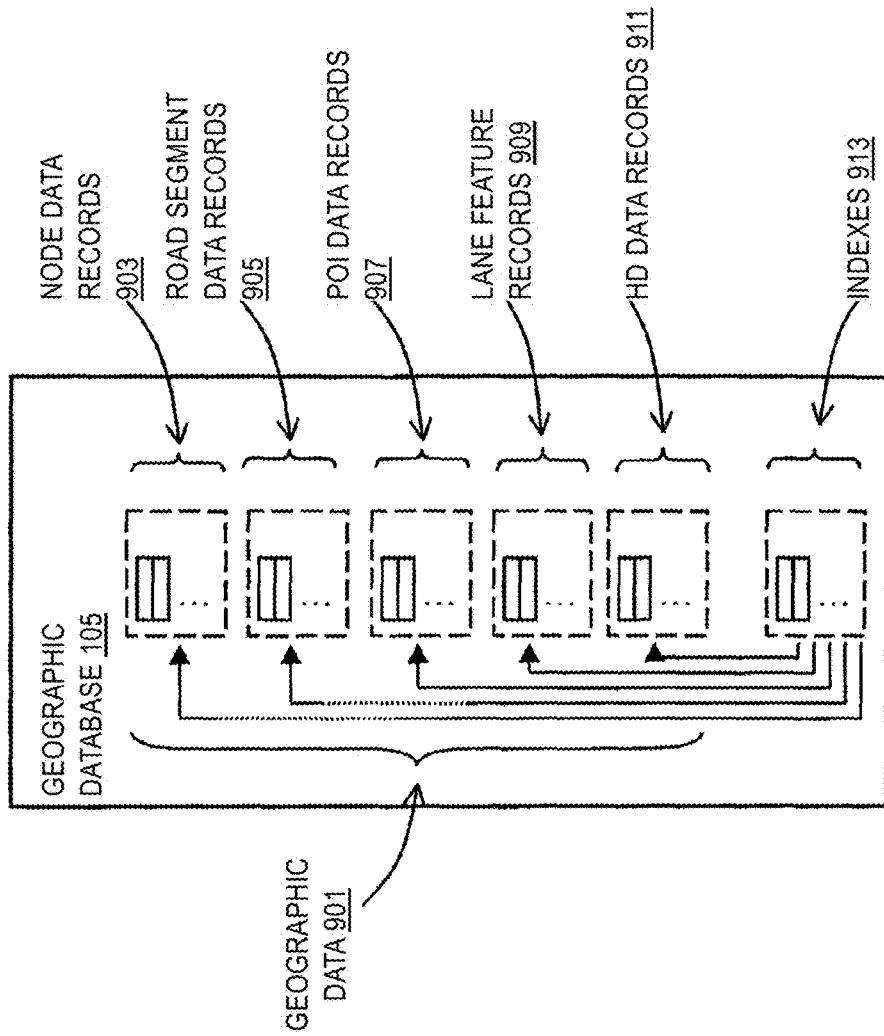


FIG. 10

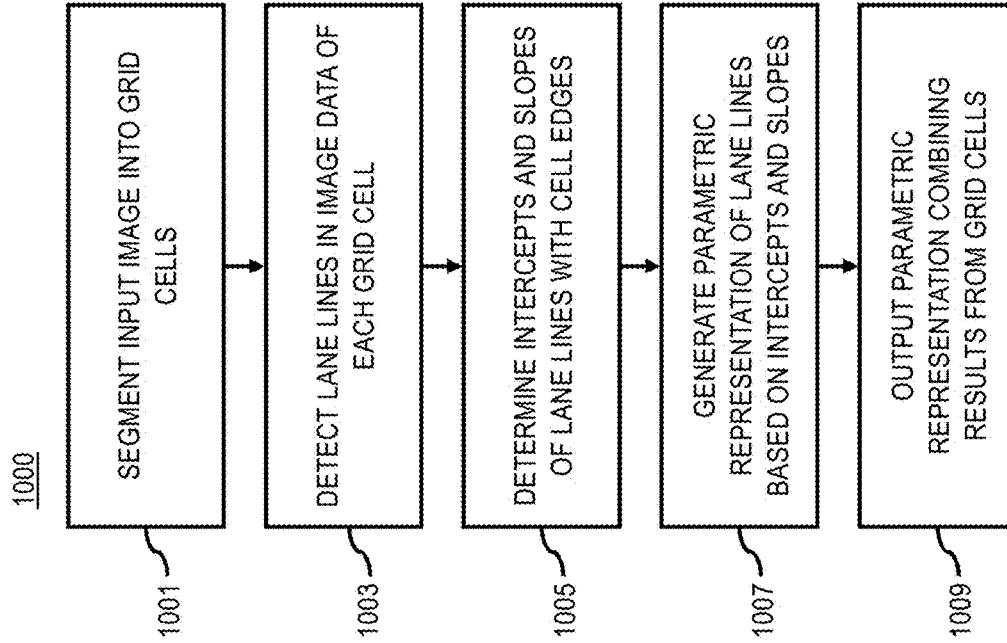


FIG. 11

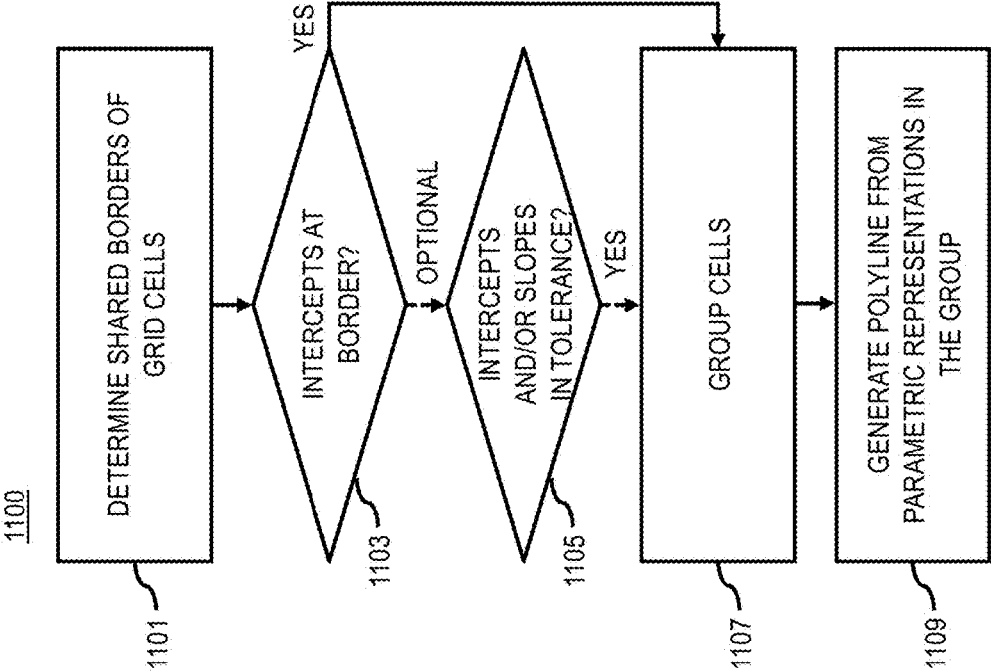


FIG. 12

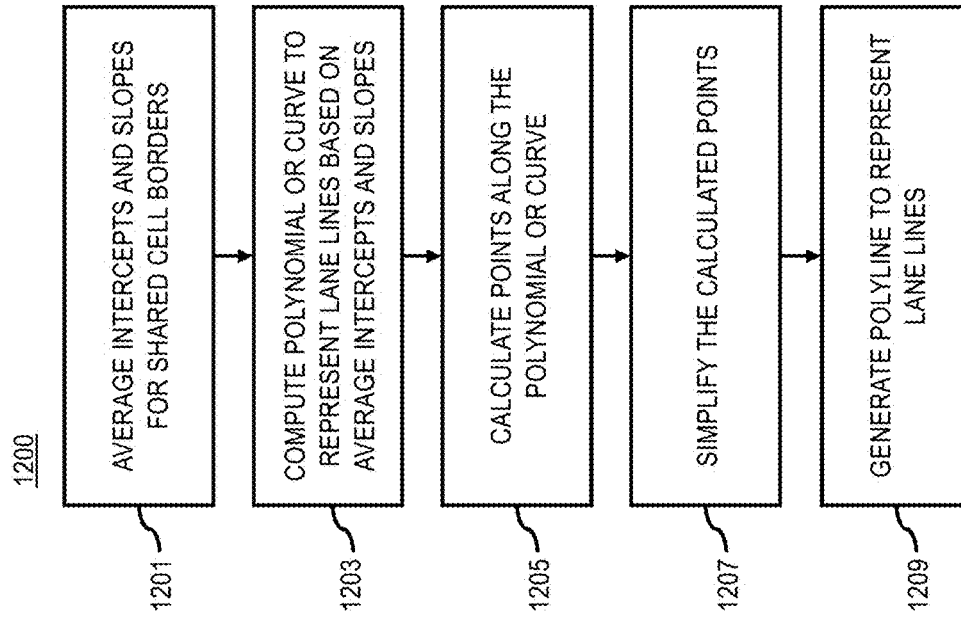


FIG. 13

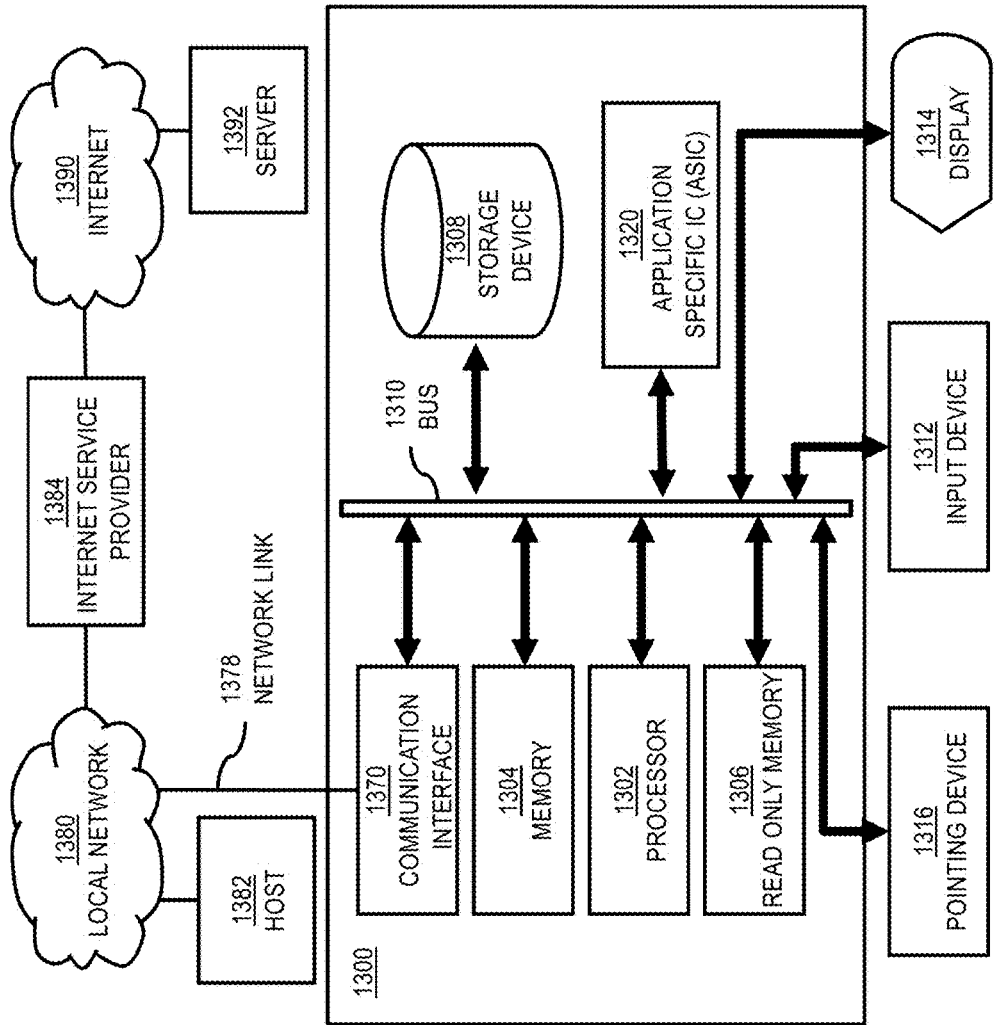


FIG. 14

1400

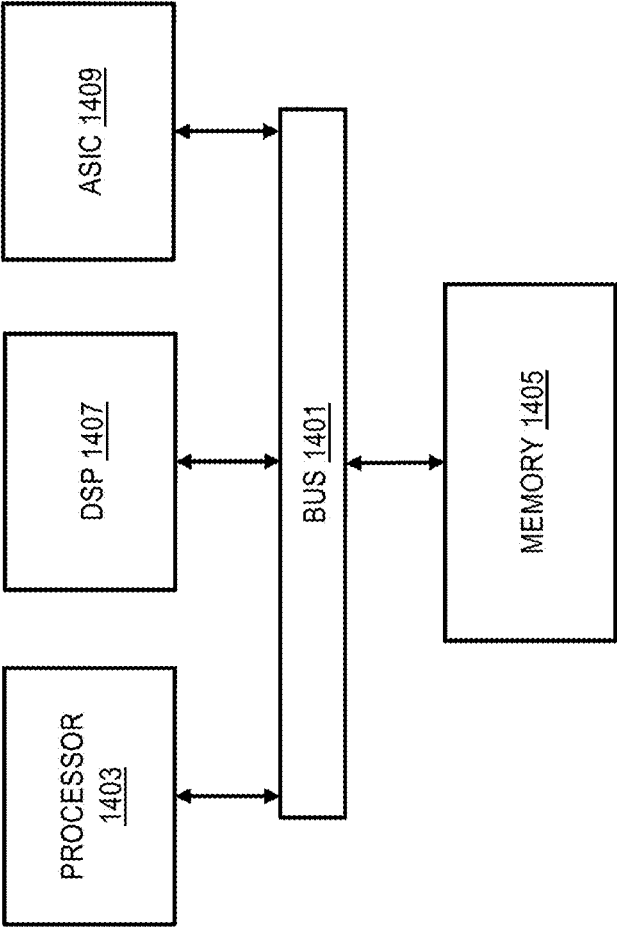
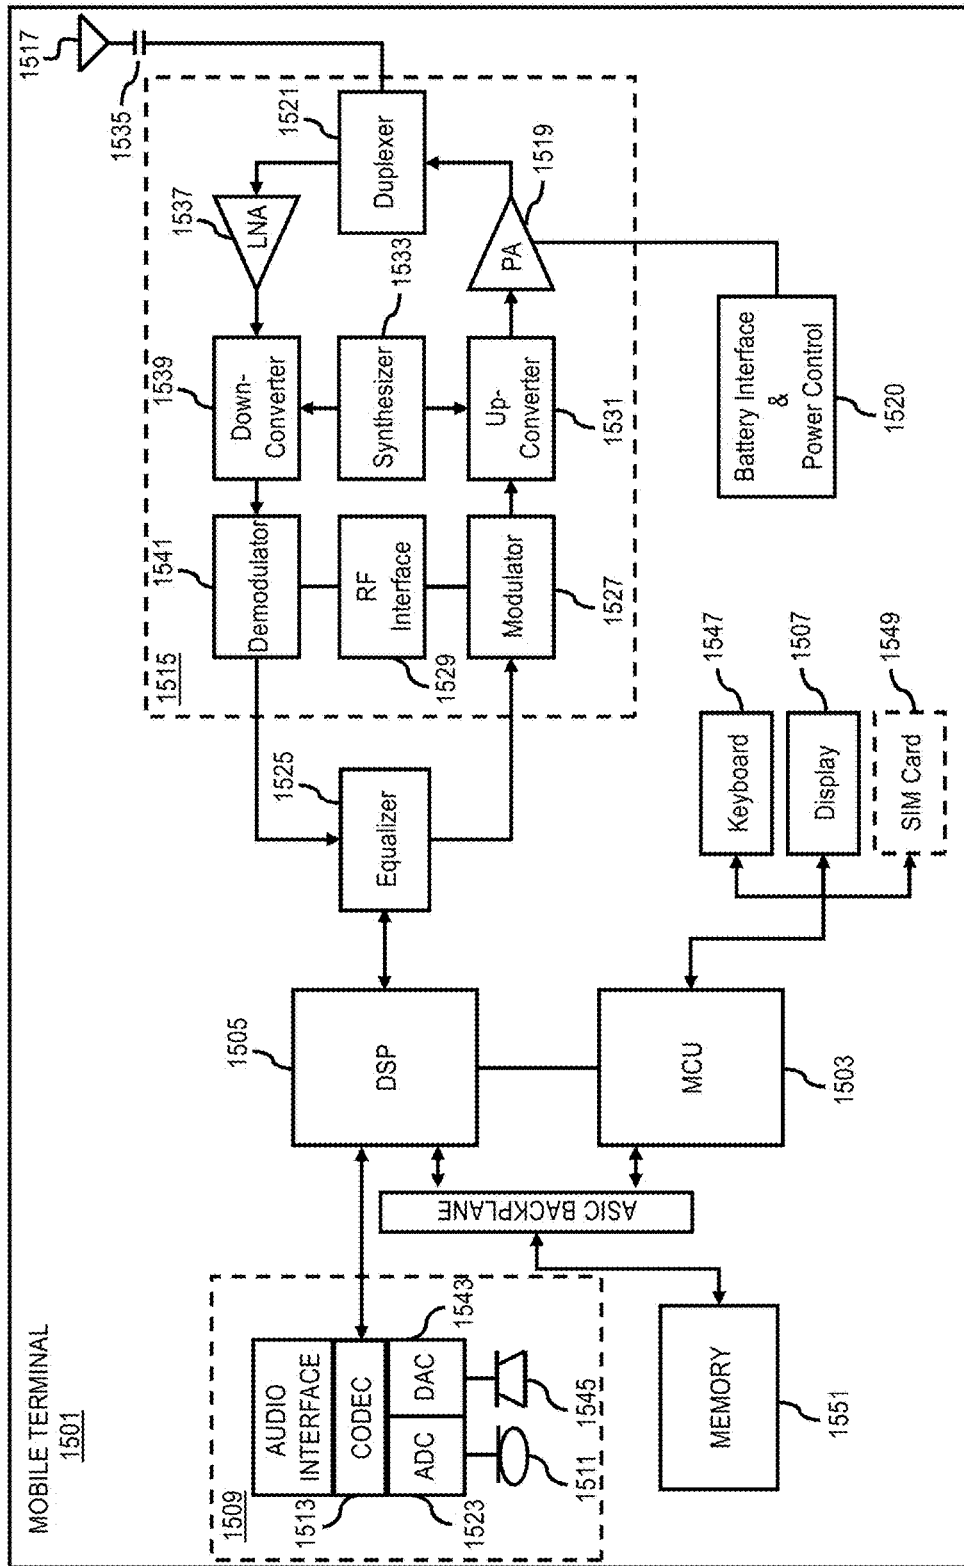


FIG. 15



## METHOD, APPARATUS, AND SYSTEM FOR VANISHING POINT/HORIZON ESTIMATION USING LANE MODELS

### BACKGROUND

**[0001]** Autonomous driving has quickly become an area of interest for vehicle manufactures and navigation and mapping service providers. One particular area of interest is the use of computer vision to enable mapping and sensing of a vehicle's environment to support autonomous or semi-autonomous operation. Advances in available computing power has enabled this mapping and sensing to approach or achieve real-time operation through, e.g., machine learning (e.g., neural networks). As a result, one application of computer vision techniques in autonomous driving is localization of the vehicle with respect to known reference marks such as lane markings and related environmental features such as the horizon or vanishing point. Accordingly, service providers face significant technical challenges when applying computer vision to estimate a horizon or vanishing point from image data.

### SOME EXAMPLE EMBODIMENTS

**[0002]** Therefore, there is a need for an approach for estimating a vanishing point or horizon from lane models automatically detected from captured images (e.g., a video capture stream from an autonomous vehicle).

**[0003]** According to one embodiment, a computer-implemented method for determining a horizon in an image depicting one or more lanes of a roadway comprises processing the image to construct one or more lane models of the one or more road lanes depicted in the image. The method also comprises extending the one or more road lanes through the image using the one or more lane models. The method further comprises determining a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line. The method further comprises designating the horizontal line as the horizon of the image.

**[0004]** According to another embodiment, an apparatus for determining a horizon in an image depicting one or more lanes of a roadway comprises at least one processor, and at least one memory including computer program code for one or more computer programs, the at least one memory and the computer program code configured to, with the at least one processor, cause, at least in part, the apparatus to process the image to construct one or more lane models of the one or more road lanes depicted in the image. The apparatus is also caused to extend the one or more road lanes through the image using the one or more lane models. The apparatus is further caused to determine a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line. The apparatus is further caused to designate the horizontal line as the horizon of the image.

**[0005]** According to another embodiment, a non-transitory computer-readable storage medium for determining a horizon in an image depicting one or more lanes of a roadway carries one or more sequences of one or more instructions which, when executed by one or more processors, cause, at least in part, an apparatus to process the image to construct one or more lane models of the one or more road lanes depicted in the image. The apparatus is also caused to

extend the one or more road lanes through the image using the one or more lane models. The apparatus is further caused to determine a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line. The apparatus is further caused to designate the horizontal line as the horizon of the image.

**[0006]** According to another embodiment, an apparatus for determining a horizon in an image depicting one or more lanes of a roadway comprises means for processing the image to construct one or more lane models of the one or more road lanes depicted in the image. The apparatus also comprises means for extending the one or more road lanes through the image using the one or more lane models. The apparatus further comprises means for determining a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line. The apparatus further comprises means for designating the horizontal line as the horizon of the image.

**[0007]** In addition, for various example embodiments of the invention, the following is applicable: a method comprising facilitating a processing of and/or processing (1) data and/or (2) information and/or (3) at least one signal, the (1) data and/or (2) information and/or (3) at least one signal based, at least in part, on (or derived at least in part from) any one or any combination of methods (or processes) disclosed in this application as relevant to any embodiment of the invention.

**[0008]** For various example embodiments of the invention, the following is also applicable: a method comprising facilitating access to at least one interface configured to allow access to at least one service, the at least one service configured to perform any one or any combination of network or service provider methods (or processes) disclosed in this application.

**[0009]** For various example embodiments of the invention, the following is also applicable: a method comprising facilitating creating and/or facilitating modifying (1) at least one device user interface element and/or (2) at least one device user interface functionality, the (1) at least one device user interface element and/or (2) at least one device user interface functionality based, at least in part, on data and/or information resulting from one or any combination of methods or processes disclosed in this application as relevant to any embodiment of the invention, and/or at least one signal resulting from one or any combination of methods (or processes) disclosed in this application as relevant to any embodiment of the invention.

**[0010]** For various example embodiments of the invention, the following is also applicable: a method comprising creating and/or modifying (1) at least one device user interface element and/or (2) at least one device user interface functionality, the (1) at least one device user interface element and/or (2) at least one device user interface functionality based at least in part on data and/or information resulting from one or any combination of methods (or processes) disclosed in this application as relevant to any embodiment of the invention, and/or at least one signal resulting from one or any combination of methods (or processes) disclosed in this application as relevant to any embodiment of the invention.

**[0011]** In various example embodiments, the methods (or processes) can be accomplished on the service provider side

or on the mobile device side or in any shared way between service provider and mobile device with actions being performed on both sides.

[0012] For various example embodiments, the following is applicable: An apparatus comprising means for performing a method of the claims.

[0013] Still other aspects, features, and advantages of the invention are readily apparent from the following detailed description, simply by illustrating a number of particular embodiments and implementations, including the best mode contemplated for carrying out the invention. The invention is also capable of other and different embodiments, and its several details can be modified in various obvious respects, all without departing from the spirit and scope of the invention. Accordingly, the drawings and description are to be regarded as illustrative in nature, and not as restrictive.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The embodiments of the invention are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings:

[0015] FIG. 1 is a diagram of a system capable of estimating a vanishing point or horizon using lane models, according to one embodiment;

[0016] FIG. 2 is a flowchart of a process for estimating a vanishing point or horizon using lane models, according to one embodiment;

[0017] FIG. 3 is a diagram of an input image of lane lines captured by a vehicle camera system, according to one embodiment;

[0018] FIGS. 4A and 4B are diagrams illustrating lane models detected from image data, according to one embodiment;

[0019] FIGS. 5A and 5B are diagrams illustrating sweeping a line across an image to estimate a vanishing point or horizon from image data, according to one embodiment;

[0020] FIG. 5C is a diagram illustrating estimating a horizon when a slope of a road varies, according to one embodiment;

[0021] FIG. 6 is a flowchart of a process for initiating mapping or navigation related functions based on an estimated vanishing point or horizon, according to one embodiment;

[0022] FIGS. 7A-7D are diagrams illustrating an example of estimating horizons over a drive to perform a consistency check, according to one embodiment;

[0023] FIGS. 8A-8D are diagrams illustrating example use cases for an estimated vanishing point or horizon, according to various embodiments;

[0024] FIG. 9 is a diagram of a geographic database, according to one embodiment;

[0025] FIG. 10 is a flowchart of a process for generating a parametric representation of lane lines detected in an input image, according to one embodiment;

[0026] FIG. 11 is a flowchart of a process for grouping grid cells based on their respective parametric representations of lane lines, according to one embodiment;

[0027] FIG. 12 is a flowchart for decoding parametric representations of lane lines into polylines, according to one embodiment;

[0028] FIG. 13 is a diagram of hardware that can be used to implement an embodiment of the invention;

[0029] FIG. 14 is a diagram of a chip set that can be used to implement an embodiment of the invention; and

[0030] FIG. 15 is a diagram of a mobile terminal (e.g., handset) that can be used to implement an embodiment of the invention.

#### DESCRIPTION OF SOME EMBODIMENTS

[0031] Examples of a method, apparatus, and computer program for estimating a vanishing point or horizon from lane models are disclosed. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the embodiments of the invention. It is apparent, however, to one skilled in the art that the embodiments of the invention may be practiced without these specific details or with an equivalent arrangement. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the embodiments of the invention.

[0032] FIG. 1 is a diagram of a system capable of estimating a vanishing point or horizon using lane models, according to one embodiment. As discussed above, autonomous driving has quickly become an area of intense interest, with recent advances in machine learning, computer vision and computing power enabling real-time mapping and sensing of a vehicle's environment. Such an understanding of the environment enables autonomous, semi-autonomous, or highly assisted driving in a vehicle (e.g., a vehicle 101) in at least two distinct ways.

[0033] First, real-time sensing of the environment provides information about potential obstacles, the behavior of others on the road, and safe, drivable areas. An understanding of where other cars are and what they might do is critical for a vehicle 101 to safely plan a route. Moreover, vehicles 101 generally must avoid both static (lamp posts, e.g.) and dynamic (cats, deer, e.g.) obstacles, and these obstacles may change or appear in real-time. More fundamentally, vehicles 101 can use a semantic understanding of what areas around them are navigable and safe for driving. Even in a situation where the world is completely mapped in high resolution, exceptions will occur in which a vehicle 101 might need to drive off the road to avoid a collision, or where a road's geometry or other map attributes like direction of travel have changed. In this case, detailed mapping may be unavailable, and the vehicle 101 has to navigate using real-time sensing of road features or obstacles using a computer vision system (e.g., a computer vision system 103).

[0034] A second application of vision techniques in autonomous driving is localization of the vehicle 101 with respect to a map of reference landmarks. Understanding one's location on a map enables planning of a route, both on fine and coarse scales. On a coarse scale, navigation maps allow vehicles 101 to know what roads to use to reach a particular destination. However, on a finer scale, maps allow vehicles 101 to know what lanes to be in and when to make lane changes. Knowing this information is important for planning an efficient and safe route, for in complicated driving situations maneuvers need to be executed in a timely fashion, and sometimes before they are visually obvious. In addition, localization with respect to a map enables the incorporation of other real-time information into route planning. Such information could include traffic, areas with unsafe driving conditions (e.g., ice, fog, potholes), and temporary road changes like construction.

[0035] With respect to lane localization and also generally with respect to autonomous driving, high accuracy and

real-time localization of vehicles **101** are needed. Traditionally, most vehicle navigation systems accomplish this localization using GPS, which generally provides a real-time location with a 95% confidence interval of 7.8 meters. However, in complicated urban environments, reflection of GPS signals can further increase this error, such that one's location may be off by as much as 30 meters. Given that the width of many lanes is 3-4 meters, this accuracy is not sufficient to properly localize a vehicle **101** (e.g., an autonomous vehicle) so that it can make safe route planning decisions. Other sensors, such as inertial measurement units (IMUs) can increase the accuracy of localization by taking into account vehicle movement, but these sensors tend to drift and still do not provide sufficient accuracy for localization.

**[0036]** In general, a localization accuracy of around 10 cm is needed for safe driving in many areas. One way to achieve this level of accuracy is to use visual odometry, in which features are detected from imagery. These features can then be matched to a database of features to determine one's location. By way of example, traditional feature-based localization that both detect features and localize against them generally rely on low-level features. However, low-level features typically used in these algorithms (e.g., Scale-Invariant Feature Transform (SIFT) or Oriented FAST and rotated BRIEF (ORB)) tend to be brittle and not persist in different environmental and lighting conditions. As a result, they often cannot be used to localize a vehicle on different days in different weather conditions. Aside from reproducibility, the ability to detect and store higher level features of different types (e.g., lane features such as lane markings, lane lines, etc.) can provide better and more accurate localization.

**[0037]** In response to these issues, the system **100** of FIG. 1 (e.g., including the computer vision system **103**) focuses on detecting high level features that have semantic meaning for human beings. One such feature that is important to autonomous driving is the detection of lane features (e.g., lane markings, lane lines, Botts' dots, reflectors, etc.) and corresponding lane models. Lane-level information is important for self-driving applications because it defines the grammar of driving. Without knowledge of lane markings, it can be difficult or impossible to determine where a vehicle **101** should drive, can drive, and what maneuvers are possible. As a result, the ability to detect lane-lines in real-time constitutes a fundamental part for the design of an autonomous vehicle **101**, and significant computer vision resources are dedicated to this task.

**[0038]** At the same time, there are other high level features such as a vanishing point or horizon that can be used to help localization or other similar navigation or mapping tasks for vehicles. As used herein, the vanishing point is the point at which receding parallel lines in a perspective view appear to converge. For example, images captured from a camera mounted on a vehicle **101** typically will generate images in this perspective view (e.g., a first person perspective of the road or environment in front of the vehicle **101**). From the perspective of real world images, a horizon is a line (e.g., an apparent horizontal line) that separates the earth from the sky. In most cases, roadways are affixed to the surface of the Earth, and the lane lines demarcating the roadways generally are uniformly parallel. Because of this, in one embodiment, the contours of roadways visible in captured images can represent the receding parallel lines of a perspective view

leading to a vanishing point, with the vanishing point located on the horizon line. Alternatively, the horizon can be considered as a two dimensional horizontal line extending from either side of the vanishing point.

**[0039]** Although the vanishing point and/or horizon can be potentially obscured in images (e.g., obscured by trees, terrain, buildings, etc.), they are features that generally can be determined from image data. However, training the computer vision system **103** to directly recognize and process images to determine a horizon or vanishing point in the same way that lane features are directly recognized and modeled (e.g., according to the embodiments of lane feature determination described herein), can significantly increase the computational burden associated with training and operating a machine learning classifier to recognize vanishing points and/or horizons in images. These additional computational demands often require service providers and manufacturers to choose which features (e.g., lane features, horizons, etc.) to target for recognition with a computer vision system in order to manage computational resources, particularly in resource constrained environments (e.g., computing environments of in-vehicle control systems).

**[0040]** To address these challenges, the system **100** of FIG. 1 introduces a capability to estimate a vanishing point or horizon in a captured image using road lane models that are automatically detected from the captured image data. In other words, the system **100** leverages detected lane models generated using the computer vision system **103** to estimate the vanishing point or horizon in the same image without using the vision techniques of the computer vision system **103** to directly recognize the vanishing point or horizon, thereby advantageously reducing computational resources that would traditionally be dedicated to estimating the vanishing point or horizon. In one embodiment, the system **100** calculates a horizon/vanishing point estimation based on a convergence of high level features (e.g., lane features) in the captured image.

**[0041]** This approach, for instance, is based on an observation that parallel lines intersect at a vanishing point in a perspective image. As noted above, images captured from the vehicle **101** as it travels along a roadway are examples of perspective images, with the viewpoint being from the camera mounting position (e.g., which can be a proxy for the driving position in the vehicle **101**). Therefore, an image captured from this position can be considered a two-dimensional (2D) projection of the three-dimensional (3D) structure or view of the roadway and/or other environment features near the vehicle **101**. Depending on the mounting position of the camera and the camera's field of view, this perspective view can be a forward view, side view (e.g., left or right side), rear view, and/or even a 360° view. This perspective view of the captured images can provide data for performing the visual odometry to localize the vehicle **101** as discussed above. In addition, parallel lines such as lane lines depicted in the images can also provide data for estimating the vanishing point or horizon according to the embodiments described herein.

**[0042]** In one embodiment, the system **100** (e.g., via a horizon estimation platform **105**) estimates a vanishing point or horizon based on detected lane models by: (a) performing a detection of road lanes in an image using a lane feature detection approach to construct lane models; (b) extending the detected lanes through the image using the lane models, e.g., towards the top of the image; and (c) performing a line

sweep, e.g., from top to bottom, of the image to find the horizontal line that gets the most lane hits (e.g., lane line intersections) within the smallest span from each other. In one embodiment, this smallest span indicates the point of convergence of the lanes (e.g., the vanishing point), and the corresponding horizontal line on which the vanishing point falls indicates the horizon. It is noted that in the embodiments described herein, the terms vanishing point and horizon are used interchangeably. Accordingly, wherever embodiments are described with respect to the vanishing point only or horizon only, the embodiments apply equally to the either vanishing point or horizon estimation.

**[0043]** In one embodiment, the system 100 can use the estimated vanishing point or horizon in a variety of situations including, but not limited to: (a) performing a consistency check for the detected lane models used to estimate the vanishing point or horizon; (b) estimating a transformation from a perspective to orthographic projection; (c) labeling and annotating map features in the image for localization; and/or (d) reducing a search area in the image for sky detection.

**[0044]** FIG. 2 is a flowchart of a process for estimating a vanishing point or horizon using lane models, according to one embodiment. In one embodiment, the horizon estimation platform 105 and/or the computer vision system 103 may perform one or more portions of the process 200 and may be implemented in, for instance, a chip set including a processor and a memory as shown in FIG. 14. As such, the horizon estimation platform 105 and/or the computer vision system 103 can provide means for accomplishing various parts of the process 200. In addition or alternatively, a services platform 107 and/or one or more services 109a-109n (also collectively referred to as services 109) may perform any combination of the steps of the process 200 in combination with the horizon estimation platform 105 and/or the computer vision system 103, or as standalone components. Although the process 200 is illustrated and described as a sequence of steps, it is contemplated that various embodiments of the process 200 may be performed in any order or combination and need not include all of the illustrated steps.

**[0045]** In step 201, the horizon estimation platform 105 processes an image depicting one or more lanes of a roadway to construct one or more lane models of the one or more road lanes depicted in the image. By way of example, lane features include, but are not limited to, lane lines, lane markings, Botts' dots, and/or any other feature indicating a travel lane of a roadway. In some cases, lane features can also include wear marks in the roadway that indicate a lane or path taken by vehicles traversing the roadway. The wear marks can include, for instance, lighter shades or groves of a road surface created by passing vehicles.

**[0046]** FIG. 3 is a diagram of an example input image 301 depicting lane lines 303a-303c captured by a vehicle camera system, according to one embodiment. In this example, the image 301 is captured in real-time by a camera system of a vehicle 101 as raster images at a predetermined pixel resolution. In one embodiment, the image 301 can be captured using cameras sensitive to visible light, infra-red, and/or any other wavelength of light. To support real-time operation, the image 301 can be part of an image stream captured at a relatively high frequency (e.g., 20 Hz, 30 Hz, or higher). Each frame of the image stream can then be processed to provide real-time detection of lane-lines. In one

embodiment, the output of a lane feature detection process include lane models representing the detected lanes or features, and associated prediction confidence values indicating a probability that the processed features of the input image are predictive of the detected lane features.

**[0047]** In one embodiment, as shown in FIG. 4A, lane models are typically represented as sets of polylines 401a-401c, in which the centerlines of the respective lanes 303a-303c are represented by piecewise-linear functions with an arbitrary number of points. In the example of FIG. 4A, the polylines 401a-401c represent each lane 303a-303a as a series of line segments (e.g., shown as dotted lines) with articulation points between the line segments indicated by circles. As previously noted, generating lane models from input images can include, but are not limited to, machine learning approaches using, for instance, neural networks to parametrically represent the lanes and compute associated confidence values (e.g., prediction confidence values).

**[0048]** FIG. 4B illustrates an example of an approach for parametrically representing lane features using a neural network that can be used to estimate a vanishing point or horizon, according to one embodiment. For example, the approach of FIG. 4B uses a more natural representation of lane lines for a neural network that is based on a grid of squares overlaid with the input image. In the example of FIG. 4B, a grid 421 segments the input image (e.g., the image 301 as shown in FIG. 3) into individual grid cells. In one embodiment, such a grid can be output by a fully convolutional neural network, which has the advantage of being computationally fast without having an excess of parameters that might lead to overfitting. For example, with respect to a neural network or other similar parallel processing system, each of the grid cells can be processed by a different neuron or processing node to more efficiently employ the available neurons or nodes and distribute the computational load for processing the entire input image. In other words, in one layer of the neural network, the scope of each neuron corresponds to the extent of the input image area within each respective grid cell. Each neuron or node can make a prediction (e.g., detection of a lane line) for each individual grid cell including computing a prediction confidence for the cell. As a result of this segmentation, the basic unit of representation then becomes each cell of the grid, in which each lane line is parametrically encoded.

**[0049]** As shown in FIG. 4B, an output parametric representation of the entire input image is superimposed on the initial input image (e.g., input image 301). In this example, line segments are drawn from the intercepts of lane lines at the appropriate angles at each cell edge. An example of one line segment is line segment 425. The remaining line segments are not labeled, but are depicted as short solid lines at each edge of a cell at which a line crosses. These series of line segments correspond for instance to respective lane lines. The output of each cell can then be combined to create lane models 423a-423c to represent the detected lane features (e.g., a polyline representation of the lane lines detected in the image). Additional description of this example approach to detecting lane features is provided below with respect to FIGS. 10-12.

**[0050]** In step 203, to find the vanishing point or horizon, the horizon estimation platform 105 extends the one or more road lanes through the image using the one or more lane models. In one embodiment, the one or more road lanes are extended towards a top of the image. For example, as shown

FIG. 5A, the input image 301 has been processed to detect lane lines 303a-303b to construct respective lane models 423a-423c as described above. In this example, the lane lines 303a-303c are visible in the image 301, but the vanishing point of those lane lines 303a-303b are obscured. For example, the lane lines 303a-303c are no longer visually distinguishable in the image 301 at a point 501. This can be due to a variety of factors such as limited camera resolution, obstructions (e.g., from trees, buildings, terrain, etc.), environmental conditions (e.g., haze, fog, darkness), camera limitations (e.g., over exposure, under exposure, out of focus, etc.), etc. Moreover, in many cases, the point 501 does not correspond to the vanishing point or horizon of the image.

[0051] Accordingly, to find the vanishing point, the horizon estimation platform 105 extends the lane lines 303a-303b beyond the point 501. In one embodiment, the extension is performed by determining the angles of the respective lane lines 303a-303b at the point 501 using the respective lane models 423a-423c. The determined angles are then used project the lines 303a-303b further through the image (e.g., to the edge of the image field). For example, when the lane models 423a-423c use cell-based parametric representations of lane lines 303a-303b as described above, the horizon estimation platform 105 can look at the last cell border of the respective lane models 423a-423c at which the respective lane lines 303a-303c cross to determine the angle of the representative polyline closest to the point 501. If any other type of lane model or representation is used, the horizon estimation platform 105 can calculate the angle at which to extend the lane lines 303a-303c using any equivalent data provided by the lane model or representation. Respective extension lines 503a-503c can be then extended from the point 501 for each lane line 303a-303c at the respectively determined angles.

[0052] In step 205, the horizon estimation platform 105 determines a horizontal line (e.g., line 505) in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line. In one embodiment, the horizon estimation platform 105 initiates a line sweep of the image from a top to a bottom of the image to determine the horizontal line 505. For example, the line sweep begins from an initial line 507 near the top of the screen to evaluate for intersections with the extended lines 503a-503c. The number of intersections between the current horizontal line and the extended lines 503a-503c are counted. The horizon estimation platform 105 also calculates an extent of the horizontal line that the intersecting lines 503a-503c cover (e.g., the left most intersecting extended line to the right most intersecting extended line). In one embodiment, each subsequent line can be similarly processed until the line meeting the criteria for a horizon is identified (e.g., a line with the maximum number intersections over a minimum extent or smallest span of the horizontal line).

[0053] FIG. 5B illustrates an example of determining the number of extended line intersections and extent of those intersections during a line sweep, according to one embodiment. More specifically, FIG. 5B illustrates three candidate horizontal lines 521a-521c of a line sweep, with each of the candidate horizontal lines 521a-521c progressing further down the input image. Each of the candidate horizontal lines 521a-521c are in an area of the input image (e.g., input image 301) where all three of the extended lines 503a-503c

intersect all three of the candidate horizontal lines 521a-521c. Accordingly, each candidate line 521a-521c has the maximum number of intersections (e.g., 3 intersections) and therefore meets the criterion for having the maximum number of intersections.

[0054] Next, the horizon estimation platform 105 calculates the extent or span of the respective candidate lines 521a-521c covered by the intersecting extended lines 503a-503c. For example, with respect to candidate horizontal line 521a, the extended lines 503a-503c cross the candidate line 521a at three intersection points 523a-523c. The span or extent 525a of the candidate line 521a from the first intersection point 523a to the third intersection point 523c is calculated. The horizon estimation platform 105 then continues the line sweep to each subsequent candidate horizontal line (e.g., candidate lines 521b and 521c) to perform similar operations, for instance, to determine an extent 525b of the candidate line 521b covered by the intersection points 527a-527c, and an extent 525c of the candidate line 521c covered by the intersection points 529a-529c.

[0055] In step 207, the horizon estimation platform 105 designates the horizontal line with the greatest number of extended lane line intersections within the shortest span or extent of the line as the horizon of the image. In the example of FIG. 5B, the candidate line 521b is designated as the horizon of the image because it has a maximum number of lane line intersections (e.g., three lane line intersections) corresponding to the smallest extent 525b among the candidate horizontal lines 521a-521b. As noted above, the smallest extent 525b represents the convergence point or vanishing point of the extended lane lines 503a-503c, and the horizontal line on which the vanishing point falls is the horizon.

[0056] In one embodiment, the estimated horizon line should not change locally except in certain cases where the slope of the road varies as shown in FIG. 5C. For example, when a vehicle 101 is driving in a horizontally flat driving position 541 the vehicle camera field of view 543 is also pointed directly at the horizon 545 resulting in a captured image 547 in which the estimated horizon 549 (e.g., determined according to the various embodiment described herein) is determined to be at approximately the middle of the image 547. However, if the vehicle 101 locally encounters a section of the roadway with a positive slope as shown in driving position 551, the vehicle camera field of view 543 is tilted upwards relative to the horizon 545 resulting in a captured image 553. Because of this upward tilt, the estimated horizon 555 determined from the captured image 553 will be relatively lower in the image than the estimated horizon 549 of image 547 even though the actual horizon 545 has not changed. Similarly, if the vehicle 101 locally encounters a section of the roadway with a negative slope as shown in driving position 557, the resulting vehicle camera field of view 543 will be tilted downwards relative to the horizon 545. This results in a captured image 559 with an estimated horizon 561 higher in the image 555 when compared to the estimated horizon 549 of image 547 even though the horizon 545 has not changed.

[0057] In one embodiment, to account for variances in the slopes of the roadway, the horizon estimation platform 105 detects a slope of the roadway, and adjusts the horizon based on the slope. The horizon estimation platform 105, for instance, can determine the slope from one or more sensors of a vehicle. For example, the vehicle 101 can be equipped

with an Inertial Measurement Unit (IMU) or other equivalent sensors to determine a slope of the vehicle **101** at the time the vehicle **101** captures an image for horizon estimation. By way of example, the IMU can sense linear and angular motion of the vehicle **101** through a combination of gyroscopes and/or accelerometers arranged according to each possible axis of motion. The IMU, for instance, collects position, heading, and attitude data (e.g., including yaw, pitch, and roll). It is contemplated that the horizon estimation platform **105** can use any process to adjust the estimated horizon to account for a sensed slope of the vehicle **101**. For example, the horizon estimation platform **105** can determine how far down an estimated horizon is to be adjusted for each degree of upward tilt or slope of the vehicle **101** from a flat horizontal position, and how far up an estimated horizon is to be adjusted for each degree of downward tilt or slope of the vehicle **101** from the flat horizontal position. In one embodiment, the magnitude of these adjustments can be based on characteristics such as, but not limited to, the vehicle **101**'s camera lens, sensors, mounting position etc.

**[0058]** In one embodiment, the adjustment of the estimated horizon normalizes the horizon estimation to a flat horizon position (e.g., no slope or tilt). For example, the estimated horizon **555** of the image **553** (e.g., captured in the driving position **551** with an upward slope) can be adjusted upwards to adjusted horizon **563** which is more consistent with the estimated horizon **549** of the flat driving position **541**. Similarly, the estimated horizon **561** of the image **559** (e.g., captured in the driving position **557** with a downward slope) can be adjusted downwards to adjusted horizon **565** which also is more consistent with the estimated horizon **549** of the flat driving position **541**.

**[0059]** FIG. **6** is a flowchart of a process for initiating mapping or navigation related functions based on an estimated vanishing point or horizon, according to one embodiment. In one embodiment, the horizon estimation platform **105** and/or the computer vision system **103** may perform one or more portions of the process **600** and may be implemented in, for instance, a chip set including a processor and a memory as shown in FIG. **14**. As such, the horizon estimation platform **105** and/or the computer vision system **103** can provide means for accomplishing various parts of the process **600**. In addition or alternatively, a services platform **107** and/or one or more services **109a-109n** (also collectively referred to as services **109**) may perform any combination of the steps of the process **600** in combination with the horizon estimation platform **105** and/or the computer vision system **103**, or as standalone components. Although the process **600** is illustrated and described as a sequence of steps, it is contemplated that various embodiments of the process **600** may be performed in any order or combination and need not include all of the illustrated steps.

**[0060]** In one embodiment, the process **600** provides optional steps or functions that can be performed using data on the quality of lane features estimated according to the various embodiments described herein. It is noted that the functions described in the process **600** are provided by way of illustration and not as limitations. Accordingly, in step **601**, the horizon estimation platform **105** determines or retrieves a vanishing point or horizon estimated from an input image using lane models according to the process **200** of FIG. **2**.

**[0061]** In step **603**, the horizon estimation platform **105** uses the horizon to determine a consistency of the one or

more lane models detected in one or more subsequent images of the one or more lanes the roadway. As described above, the estimated horizon should be locally consistent (e.g., consistent within a local or specified geographic area, or locally with respect to a given trip of the vehicle **101**). Exceptions to this can include variances in roadway slope or other similar factor. As discussed above, in one embodiment, the horizon estimation platform **105** can adjust the estimated horizons to account or normalize for these variances.

**[0062]** In one embodiment, to perform a consistency check on lane models, the horizon estimation platform **105** can capture or retrieve a sequence of images covering the times and/or locations for which the check is requested. The horizon estimation platform **105** then processes each image to determine respective estimated horizons, accounting for slope variances as needed. The horizon estimation platform **105** can then determine whether each estimated horizon is equivalent (e.g., statistically equivalent using any known outlier or equivalent test). In one embodiment, any outlying or inconsistent estimated horizon and/or corresponding image used to generate the estimated horizon can be flagged for additional evaluation.

**[0063]** FIGS. **7A-7D** are diagrams illustrating an example of estimating horizons over a drive to perform a consistency check, according to one embodiment. In the example of FIGS. **7A-7B**, a vehicle **101** is driving along a route, and a graph **701** illustrates the estimated horizon determined over time as the vehicle **101** drives. Each dot depicted in the graph **701** represents a separate input image (e.g., a video frame of a sequence of images captured from a camera mounted on the vehicle **101**) processed to determine an estimated horizon according to the various embodiments described herein. The average of the estimated horizons over this drive is indicated by a line **703** (e.g., at vertical pixel **300**). In this example, the unit of the estimated horizon line is given as a vertical pixel number of the input image corresponding to a horizontal line of pixels representing the estimated horizon. This notation is provided by way of illustration and is not intended as a limitation. Accordingly, it is contemplated that the horizon estimation platform **105** can use any notation to indicate the estimated horizon.

**[0064]** Each of the FIGS. **7A-7D** illustrates estimating a horizon from a respective image taken at a different point in time of drive illustrated in the graph **701**. For example, FIG. **7A** illustrates a time **705** in the drive at which the vehicle **101** captured an image **707**. The image **707** provides an example of an image in which lane lines are clearly visible, thereby resulting in the detection of high quality lane models **707a-707d** from which an estimated horizon **709** is computed. For example, the lane models **707a-707d** are characterized as high quality for horizon estimation because lane lines from which the models are estimated are visible over a large portion of the image, so that more grid cells can be used to parametrically represent the lines as the lane models **707a-707d**. Because of this, the lane models **707a-707d** can be more likely to accurately represent the lane lines in the image **705**, resulting in a convergence point **711** of the extended lane models **707a-707d** that are tightly grouped over a smaller extent of the horizontal line of the estimated horizon **709**. As a result, the estimated horizon **709** is consistent with the average estimated horizon **703**, and therefore the horizon estimation platform **105** can designate the lane models **707a-707d** as consistent.

[0065] FIG. 7B illustrates a time 721 in the drive at which the vehicle 101 captured an image 723. In this case, the vehicle 101 is in a tunnel, leading to an overexposure of the area corresponding to the exit of the tunnel. The overexposed area of the image 723 further obscures the horizon. However, the embodiments of horizon estimation according described herein can continue to work very well in an overexposed image as long as limited segments of lane lines are detected and can be extended through the image. In this example, despite the overexposed area, significant segments of at least three lane lines remain visible in the image 723 from which lane models 725a-725c can be detected. The lane models 725a-725c are then extended through the overexposed areas to find a convergence point 727 and determine the estimated horizon 729. Despite the overexposed area, the estimated horizon 729 remains consistent with the average estimated horizon 703, and therefore the horizon estimation platform 105 can designate the lane models 725a-725c as consistent.

[0066] FIG. 7C illustrates a time 741 in the drive at which the vehicle 101 captured an image 743. In this example, the horizon estimation platform 105 detects only one lane model 745 and extends it through the image 743. However, because only one lane model 745 was detected, the horizon estimation platform 105 cannot determine a convergence point or an estimated horizon. As a result, no estimated horizon is reported for the time 741, and a horizon estimate of 0 is recorded in the graph 701. When performing a consistency check of the lane model 745 for the time 741, the zero result can be noted as an outlier. The lane model 745 can then be marked as inconsistent. Based on this designation, the horizon estimation platform 105, for instance, can ignore the image 741 and the corresponding lane model 745, mark the lane model 745 as potentially low quality, designate the road segment as needing lane maintenance (e.g., poor quality lane markings, etc.), and/or take any other equivalent action.

[0067] FIG. 7D illustrates a time 761 in the drive at which the vehicle 101 captured an image 763. In this example, the image 763 captures a portion of the roadway where one subset of lanes is going one direction and another subset is going in another direction. This can occur for instance, at a highway exit or split as shown, where one set of lanes correspond to the main road, and another set of lanes correspond to the exit ramp. This can be problematic for the horizon estimation platform 105 because there could potentially be multiple convergence points from the lanes going different directions. For example, detected lane models 765a-765c correspond to the main road going one direction and converging at point 767, and detected lane models 769a-769b correspond to an exit ramp going another direction and converging at point 771. If only one lane model is detected for each direction, it may not be possible to determine the convergence points (e.g., similar to case of FIG. 7C).

[0068] In one embodiment, if there are multiple convergence points, the horizon estimation platform 105 can cluster each set of lanes according to their convergence points to determine estimated horizons for each convergence point independent using the embodiment described above. The horizon estimation platform 105 can then further process the resulting multiple estimated horizons to determine a consensus horizon 773 for the image 763 (e.g., determine an average as the consensus estimated horizon, select the estimated horizon generated from the highest quality lane

models, etc.). In cases where the multiple convergence points fall on the same or close to the same (e.g., within a threshold value) horizontal line, the resulting consensus estimated horizon may be consistent with the average estimated horizon 703. However, in the example of FIG. 7D, the main road and the exit ramp differ in slope (e.g., the exit ramp slopes to a higher level) so that the convergence points 767 and 771 occur on significantly different horizontal lines. This leads to an estimated horizon 773 (e.g., 200 pixels) that is an outlier and inconsistent with the average estimated horizon 703. Accordingly, the horizon estimation platform 105 can flag the image 763 and/or the lane models 765a-765c and 767a-767b as inconsistent for further evaluation or take other actions to reduce the impact of the outlier (e.g., actions as discussed with respect to FIG. 7C).

[0069] FIG. 8A is diagram illustrating an example vehicle navigation system that can initiate functions based on a horizon-based lane model inconsistency check, according to one embodiment. As shown in FIG. 8A, a vehicle 101 is traveling along a roadway 801 in fully autonomous mode using a navigation/vehicle control system 803 configured to continuously perform a consistency of the detected lane models based on estimated horizons as described above. In this example, the system 803 (e.g., interacting with the horizon estimation platform 105) detects that the vehicle 101 is approaching a segment of the roadway 801 wherein estimated horizons are inconsistent above a predetermined threshold (e.g., a statistically significant outlier). In one embodiment, the system 803 is configured to interpret these horizon inconsistencies as indicating a potential problems. These problems can include, but are not limited to, problems with then computer vision system 103, problems with the roadway that would make accurate localization less likely, etc. Accordingly, the system 803 presents an alert message 805 to the driver indicating "Road Alert! Potential horizon inconsistency in lane models" and providing the driver with an option to exit autonomous mode to take manual control of the vehicle 101.

[0070] Returning to the process 600 of illustrating use cases for the estimated horizon, in step 605, the horizon estimation platform 105 estimates a transformation of a perspective representation to an orthographic projection of the one or more road lanes based on the horizon and the one or more lane models. In other words, as shown in FIG. 8, the horizon estimation platform 105 can use the horizon 821 estimated from the input image 823 to convert the perspective view of the image 823 into some other orthographic projection of the same view to simulate a different apparent camera position. For example, the input image generally provides a first-person view from a driving position of the vehicle looking outwards towards upcoming sections of the roadway as shown in the image 601. In one embodiment, the horizon estimation platform 105 can use the estimated horizon 821 determined from the image 823 as a reference point from which to render or transform the perspective view into an orthographic projection 825 of the same scene. In other words, the estimated horizon 821 can be used to help localize the vehicle 101 on the roadway depicted in the image 823. This localization is then used to render the orthographic projection 825 which simulates a picture taken not from the first person perspective of the vehicle 101 as shown in the image 823, but from a simulated camera

position from the outside front and to the left of the vehicle **101**, while on the same road at the same time as the original image **823**.

[0071] In step **607**, the horizon estimation platform **105** labels or annotates one or more map features depicted in the image based on the horizon. As shown in FIG. **8C**, the horizon estimation platform **105** can determine an estimated horizon **821** from the image **843** according to the embodiments described herein. The estimated horizon **821** can be used as a reference point for identifying other features (e.g., a sign **845**, or other roadside objects) in the image **843**. By way of example, the geographic database **111** that supports mapping and navigation functions of the horizon estimation platform **105** may be a high definition (HD) maps that provide centimeter-level or better positioning information for roadways include signs and objects along the roadway (e.g., to support visual odometry for localization). The horizon estimation platform **105** can localize the any detected objects in the image **843** (e.g., the sign **845**) by localizing the sign **845** based, at least in part, on a distance **847** to the estimated horizon, and then identifying the localized sign **845** by querying the geographic database for the HD location information for a matching sign in the geographic database **111**. The stored information for the matching sign can then be used to label or annotate the sign **845** as depicted in the image **843**. If no such sign is present in the geographic database **111**, the detected sign **845** can be marked for possible inclusion in or update to the geographic database **111** as a new data record.

[0072] In one embodiment, the localization of the vehicle **101** is then further based on the labeling or the annotating of the sign **845**. For example, the precise location of the sign **845** can be determined from the geographic database **111** and/or the estimated horizon **841**. The vehicle **101**'s position with respect to the sign **845** can then be used to localize the vehicle **101** by knowing the true position of the sign **845** and the relative position of the vehicle **101** to the sign **845** through, e.g., visual odometry.

[0073] In step **609**, the horizon estimation platform **105** selects a portion of the image for performing a sky detection based on the horizon. In other words, as shown in FIG. **8D**, the horizon estimation platform **105** can use a horizon **861** estimated from the lane models detected in the image **863** to segment the image **863** into a sky portion **865** (e.g., the portion of the image **863** above the estimated horizon **861**) and a non-sky portion **867** (e.g., the portion of the image **863** below the estimated horizon **861**). The horizon estimation platform **105** can then initiate any sky detection process on the just the sky portion **865** to advantageously reduce the computation resources associated with performing sky detection on the entire image. In this way, the amount of image data to process for sky detection is reduced by focusing on those areas of the image **863** most likely to include depictions of the sky.

[0074] Returning to FIG. **1**, as shown, the system **100** includes the horizon estimation platform **105** for estimating a horizon in an image using lane models according the various embodiments described herein. In addition, the system **100** includes the computer vision system **103** configured to detect lane lines in an input image and generate lane models for processing by the horizon estimation platform **105**, according to the various embodiments described herein. In one embodiment, the computer vision system **103** includes a neural network or other machine learning/parallel

processing system to automatically detect features such as lane lines in image data to support localization of, e.g., a vehicle **101** within the sensed environment. In one embodiment, the neural network of the computer vision system **103** is a traditional convolutional neural network which consists of multiple layers of collections of one or more neurons (e.g., processing nodes of the neural network) which are configured to process a portion of an input image. In one embodiment, the receptive fields of these collections of neurons (e.g., a receptive layer) can be configured to correspond to the area of an input image delineated by a respective a grid cell generated as described above.

[0075] In one embodiment, the computer vision system **103** also has connectivity or access to a geographic database **111** which representations of mapped geographic features to facilitate visual odometry to increase localization accuracy. The geographic database **111** can also store parametric representations of lane lines and other similar features and/or related data generated or used to encode or decode parametric representations of lane lines (e.g., lane models) according to the various embodiments described herein.

[0076] In one embodiment, the computer vision system **103** has connectivity over a communication network **113** to the services platform **107** that provides one or more services **109**. By way of example, the services **109** may be third party services and include mapping services, navigation services, travel planning services, notification services, social networking services, content (e.g., audio, video, images, etc.) provisioning services, application services, storage services, contextual information determination services, location based services, information based services (e.g., weather, news, etc.), etc. In one embodiment, the services **109** uses the output of the horizon estimation platform **105** (e.g., estimated lane feature quality) and/or of the computer vision system **103** (e.g., detected lane features) to localize the vehicle **101** or a user equipment **115** (e.g., a portable navigation device, smartphone, portable computer, tablet, etc.) to provide services **109** such as navigation, mapping, other location-based services, etc.

[0077] In one embodiment, the horizon estimation platform **105** may be a platform with multiple interconnected components. The horizon estimation platform **105** may include multiple servers, intelligent networking devices, computing devices, components and corresponding software for providing parametric representations of lane lines. In addition, it is noted that the horizon estimation platform **105** and/or the computer vision system **103** may be a separate entity of the system **100**, a part of the one or more services **109**, a part of the services platform **107**, or included within the UE **115** and/or vehicle **101**.

[0078] In one embodiment, content providers **117a-117m** (collectively referred to as content providers **117**) may provide content or data (e.g., including geographic data, parametric representations of mapped features, etc.) to the geographic database **111**, the horizon estimation platform **105**, the computer vision system **103**, the services platform **107**, the services **109**, the UE **115**, the vehicle **101**, and/or an application **119** executing on the UE **115**. The content provided may be any type of content, such as map content, textual content, audio content, video content, image content, etc. In one embodiment, the content providers **117** may provide content that may aid in the detecting and classifying of lane lines and/or other features in image data, and estimating the quality of the detected features. In one

embodiment, the content providers **117** may also store content associated with the geographic database **111**, horizon estimation platform **105**, computer vision system **103**, services platform **107**, services **109**, UE **115**, and/or vehicle **101**. In another embodiment, the content providers **117** may manage access to a central repository of data, and offer a consistent, standard interface to data, such as a repository of the geographic database **111**.

**[0079]** In one embodiment, the UE **115** and/or vehicle **101** may execute a software application **119** to collect, encode, and/or decode lane feature data detected in image data to estimate the quality of the lane features according to the embodiments described herein. By way of example, the application **119** may also be any type of application that is executable on the UE **115** and/or vehicle **101**, such as autonomous driving applications, mapping applications, location-based service applications, navigation applications, content provisioning services, camera/imaging application, media player applications, social networking applications, calendar applications, and the like. In one embodiment, the application **119** may act as a client for the horizon estimation platform **105** and/or the computer vision system **103** and perform one or more functions associated with estimating the quality of lane features alone or in combination with the horizon estimation platform **105**.

**[0080]** By way of example, the UE **115** is any type of embedded system, mobile terminal, fixed terminal, or portable terminal including a built-in navigation system, a personal navigation device, mobile handset, station, unit, device, multimedia computer, multimedia tablet, Internet node, communicator, desktop computer, laptop computer, notebook computer, netbook computer, tablet computer, personal communication system (PCS) device, personal digital assistants (PDAs), audio/video player, digital camera/camcorder, positioning device, fitness device, television receiver, radio broadcast receiver, electronic book device, game device, or any combination thereof, including the accessories and peripherals of these devices, or any combination thereof. It is also contemplated that the UE **115** can support any type of interface to the user (such as “wearable” circuitry, etc.). In one embodiment, the UE **115** may be associated with the vehicle **101** or be a component part of the vehicle **101**.

**[0081]** In one embodiment, the UE **115** and/or vehicle **101** are configured with various sensors for generating or collecting environmental image data (e.g., for processing the horizon estimation platform **105** and/or computer vision system **103**), related geographic data, etc. In one embodiment, the sensed data represent sensor data associated with a geographic location or coordinates at which the sensor data was collected. By way of example, the sensors may include a global positioning sensor for gathering location data (e.g., GPS), IMUs, a network detection sensor for detecting wireless signals or receivers for different short-range communications (e.g., Bluetooth, Wi-Fi, Li-Fi, near field communication (NFC) etc.), temporal information sensors, a camera/imaging sensor for gathering image data (e.g., the camera sensors may automatically capture road sign information, images of road obstructions, etc. for analysis), an audio recorder for gathering audio data, velocity sensors mounted on steering wheels of the vehicles, switch sensors for determining whether one or more vehicle switches are engaged, and the like.

**[0082]** Other examples of sensors of the UE **115** and/or vehicle **101** may include light sensors, orientation sensors augmented with height sensors and acceleration sensor (e.g., an accelerometer can measure acceleration and can be used to determine orientation of the vehicle), tilt sensors to detect the degree of incline or decline (e.g., slope) of the vehicle along a path of travel, moisture sensors, pressure sensors, etc. In a further example embodiment, sensors about the perimeter of the UE **115** and/or vehicle **101** may detect the relative distance of the vehicle from a lane or roadway, the presence of other vehicles, pedestrians, traffic lights, potholes and any other objects, or a combination thereof. In one scenario, the sensors may detect weather data, traffic information, or a combination thereof. In one embodiment, the UE **115** and/or vehicle **101** may include GPS or other satellite-based receivers to obtain geographic coordinates from satellites **121** for determining current location and time. Further, the location can be determined by a triangulation system such as A-GPS, Cell of Origin, or other location extrapolation technologies. In yet another embodiment, the sensors can determine the status of various control elements of the car, such as activation of wipers, use of a brake pedal, use of an acceleration pedal, angle of the steering wheel, activation of hazard lights, activation of head lights, etc.

**[0083]** In one embodiment, the communication network **113** of system **100** includes one or more networks such as a data network, a wireless network, a telephony network, or any combination thereof. It is contemplated that the data network may be any local area network (LAN), metropolitan area network (MAN), wide area network (WAN), a public data network (e.g., the Internet), short range wireless network, or any other suitable packet-switched network, such as a commercially owned, proprietary packet-switched network, e.g., a proprietary cable or fiber-optic network, and the like, or any combination thereof. In addition, the wireless network may be, for example, a cellular network and may employ various technologies including enhanced data rates for global evolution (EDGE), general packet radio service (GPRS), global system for mobile communications (GSM), Internet protocol multimedia subsystem (IMS), universal mobile telecommunications system (UMTS), etc., as well as any other suitable wireless medium, e.g., worldwide interoperability for microwave access (WiMAX), Long Term Evolution (LTE) networks, code division multiple access (CDMA), wideband code division multiple access (WCDMA), wireless fidelity (Wi-Fi), wireless LAN (WLAN), Bluetooth®, Internet Protocol (IP) data casting, satellite, mobile ad-hoc network (MANET), and the like, or any combination thereof.

**[0084]** By way of example, the horizon estimation platform **105**, computer vision system **103**, services platform **107**, services **109**, UE **115**, vehicle **101**, and/or content providers **117** communicate with each other and other components of the system **100** using well known, new or still developing protocols. In this context, a protocol includes a set of rules defining how the network nodes within the communication network **113** interact with each other based on information sent over the communication links. The protocols are effective at different layers of operation within each node, from generating and receiving physical signals of various types, to selecting a link for transferring those signals, to the format of information indicated by those signals, to identifying which software application executing

on a computer system sends or receives the information. The conceptually different layers of protocols for exchanging information over a network are described in the Open Systems Interconnection (OSI) Reference Model.

**[0085]** Communications between the network nodes are typically effected by exchanging discrete packets of data. Each packet typically comprises (1) header information associated with a particular protocol, and (2) payload information that follows the header information and contains information that may be processed independently of that particular protocol. In some protocols, the packet includes (3) trailer information following the payload and indicating the end of the payload information. The header includes information such as the source of the packet, its destination, the length of the payload, and other properties used by the protocol. Often, the data in the payload for the particular protocol includes a header and payload for a different protocol associated with a different, higher layer of the OSI Reference Model. The header for a particular protocol typically indicates a type for the next protocol contained in its payload. The higher layer protocol is said to be encapsulated in the lower layer protocol. The headers included in a packet traversing multiple heterogeneous networks, such as the Internet, typically include a physical (layer 1) header, a data-link (layer 2) header, an internetwork (layer 3) header and a transport (layer 4) header, and various application (layer 5, layer 6 and layer 7) headers as defined by the OSI Reference Model.

**[0086]** FIG. 9 is a diagram of a geographic database, according to one embodiment. In one embodiment, the geographic database **111** includes geographic data **901** used for (or configured to be compiled to be used for) mapping and/or navigation-related services, such as for visual odometry based on the parametric representation of lanes include, e.g., encoding and/or decoding parametric representations into lane lines. In one embodiment, the geographic database **111** include high definition (HD) mapping data that provide centimeter-level or better accuracy of map features. For example, the geographic database **111** can be based on Light Detection and Ranging (LiDAR) or equivalent technology to collect billions of 3D points and model road surfaces and other map features down to the number lanes and their widths. In one embodiment, the HD mapping data (e.g., HD data records **911**) capture and store details such as the slope and curvature of the road, lane markings, roadside objects such as sign posts, including what the signage denotes. By way of example, the HD mapping data enable highly automated vehicles to precisely localize themselves on the road.

**[0087]** In one embodiment, geographic features (e.g., two-dimensional or three-dimensional features) are represented using polygons (e.g., two-dimensional features) or polygon extrusions (e.g., three-dimensional features). For example, the edges of the polygons correspond to the boundaries or edges of the respective geographic feature. In the case of a building, a two-dimensional polygon can be used to represent a footprint of the building, and a three-dimensional polygon extrusion can be used to represent the three-dimensional surfaces of the building. Accordingly, the terms polygons and polygon extrusions as used herein can be used interchangeably.

**[0088]** In one embodiment, the following terminology applies to the representation of geographic features in the geographic database **111**.

**[0089]** “Node”—A point that terminates a link.

**[0090]** “Line segment”—A straight line connecting two points.

**[0091]** “Link” (or “edge”)—A contiguous, non-branching string of one or more line segments terminating in a node at each end.

**[0092]** “Shape point”—A point along a link between two nodes (e.g., used to alter a shape of the link without defining new nodes).

**[0093]** “Oriented link”—A link that has a starting node (referred to as the “reference node”) and an ending node (referred to as the “non reference node”).

**[0094]** “Simple polygon”—An interior area of an outer boundary formed by a string of oriented links that begins and ends in one node. In one embodiment, a simple polygon does not cross itself.

**[0095]** “Polygon”—An area bounded by an outer boundary and none or at least one interior boundary (e.g., a hole or island). In one embodiment, a polygon is constructed from one outer simple polygon and none or at least one inner simple polygon. A polygon is simple if it just consists of one simple polygon, or complex if it has at least one inner simple polygon.

**[0096]** In one embodiment, the geographic database **111** follows certain conventions. For example, links do not cross themselves and do not cross each other except at a node. Also, there are no duplicated shape points, nodes, or links. Two links that connect each other have a common node. In the geographic database **111**, overlapping geographic features are represented by overlapping polygons. When polygons overlap, the boundary of one polygon crosses the boundary of the other polygon. In the geographic database **111**, the location at which the boundary of one polygon intersects they boundary of another polygon is represented by a node. In one embodiment, a node may be used to represent other locations along the boundary of a polygon than a location at which the boundary of the polygon intersects the boundary of another polygon. In one embodiment, a shape point is not used to represent a point at which the boundary of a polygon intersects the boundary of another polygon.

**[0097]** As shown, the geographic database **111** includes node data records **903**, road segment or link data records **905**, POI data records **907**, lane feature records **909**, HD mapping data records **911**, and indexes **913**, for example. More, fewer or different data records can be provided. In one embodiment, additional data records (not shown) can include cartographic (“carto”) data records, routing data, and maneuver data. In one embodiment, the indexes **913** may improve the speed of data retrieval operations in the geographic database **111**. In one embodiment, the indexes **913** may be used to quickly locate data without having to search every row in the geographic database **111** every time it is accessed. For example, in one embodiment, the indexes **913** can be a spatial index of the polygon points associated with stored feature polygons.

**[0098]** In exemplary embodiments, the road segment data records **905** are links or segments representing roads, streets, or paths, as can be used in the calculated route or recorded route information for determination of one or more personalized routes. The node data records **903** are end points corresponding to the respective links or segments of the road segment data records **905**. The road link data records **905** and the node data records **903** represent a road network, such

as used by vehicles, cars, and/or other entities. Alternatively, the geographic database **111** can contain path segment and node data records or other data that represent pedestrian paths or areas in addition to or instead of the vehicle road record data, for example.

**[0099]** The road/link segments and nodes can be associated with attributes, such as geographic coordinates, street names, address ranges, speed limits, turn restrictions at intersections, and other navigation related attributes, as well as POIs, such as gasoline stations, hotels, restaurants, museums, stadiums, offices, automobile dealerships, auto repair shops, buildings, stores, parks, etc. The geographic database **111** can include data about the POIs and their respective locations in the POI data records **907**. The geographic database **111** can also include data about places, such as cities, towns, or other communities, and other geographic features, such as bodies of water, mountain ranges, etc. Such place or feature data can be part of the POI data records **307** or can be associated with POIs or POI data records **907** (such as a data point used for displaying or representing a position of a city).

**[0100]** In one embodiment, the geographic database **111** can also include lane feature records **909** for storing the lane lines (e.g., parametric representations of lane lines) detected from input image data according to the various embodiments described herein. In one embodiment, the geographic database **111** can also store the confidence values and the estimated quality of the detected lane features. By way of example, the lane feature records **909** can be associated with one or more of the node records **903**, road segment records **905**, and/or POI data records **907** to support localization or visual odometry based on the features stored therein and the corresponding estimated quality of the features. In this way, the parametric representation records **909** can also be associated with the characteristics or metadata of the corresponding record **903**, **905**, and/or **907**.

**[0101]** In one embodiment, as discussed above, the HD mapping data records **911** model road surfaces and other map features to centimeter-level or better accuracy. The HD mapping data records **911** also include lane models that provide the precise lane geometry with lane boundaries, as well as rich attributes of the lane models. These rich attributes include, but are not limited to, lane traversal information, lane types, lane marking types, lane level speed limit information, and/or the like. In one embodiment, the HD mapping data records **911** are divided into spatial partitions of varying sizes to provide HD mapping data to vehicles **101** and other end user devices with near real-time speed without overloading the available resources of the vehicles **101** and/or devices (e.g., computational, memory, bandwidth, etc. resources).

**[0102]** In one embodiment, the HD mapping data records **911** are created from high-resolution 3D mesh or point-cloud data generated, for instance, from LiDAR-equipped vehicles. The 3D mesh or point-cloud data are processed to create 3D representations of a street or geographic environment at centimeter-level accuracy for storage in the HD mapping data records **911**.

**[0103]** In one embodiment, the HD mapping data records **911** also include real-time sensor data collected from probe vehicles in the field. The real-time sensor data, for instance, integrates real-time traffic information, weather, and road conditions (e.g., potholes, road friction, road wear, etc.) with highly detailed 3D representations of street and geographic

features to provide precise real-time also at centimeter-level accuracy. Other sensor data can include vehicle telemetry or operational data such as windshield wiper activation state, braking state, steering angle, accelerator position, and/or the like.

**[0104]** In one embodiment, the geographic database **111** can be maintained by the content provider **117** in association with the services platform **107** (e.g., a map developer). The map developer can collect geographic data to generate and enhance the geographic database **111**. There can be different ways used by the map developer to collect data. These ways can include obtaining data from other sources, such as municipalities or respective geographic authorities. In addition, the map developer can employ field personnel to travel by vehicle (e.g., vehicle **101** and/or UE **115**) along roads throughout the geographic region to observe features and/or record information about them, for example. Also, remote sensing, such as aerial or satellite photography, can be used.

**[0105]** The geographic database **111** can be a master geographic database stored in a format that facilitates updating, maintenance, and development. For example, the master geographic database or data in the master geographic database can be in an Oracle spatial format or other spatial format, such as for development or production purposes. The Oracle spatial format or development/production database can be compiled into a delivery format, such as a geographic data files (GDF) format. The data in the production and/or delivery formats can be compiled or further compiled to form geographic database products or databases, which can be used in end user navigation devices or systems.

**[0106]** For example, geographic data is compiled (such as into a platform specification format (PSF) format) to organize and/or configure the data for performing navigation-related functions and/or services, such as route calculation, route guidance, map display, speed calculation, distance and travel time functions, and other functions, by a navigation device, such as by a vehicle **101** or UE **115**, for example. The navigation-related functions can correspond to vehicle navigation, pedestrian navigation, or other types of navigation. The compilation to produce the end user databases can be performed by a party or entity separate from the map developer. For example, a customer of the map developer, such as a navigation device developer or other end user device developer, can perform compilation on a received geographic database in a delivery format to produce one or more compiled navigation databases.

**[0107]** FIG. **10** is a flowchart of a process for generating a parametric representation of lane lines detected in an input image, according to one embodiment. In one embodiment, the computer vision system **103** and/or the horizon estimation platform **105** may perform one or more portions of the process **1000** and may be implemented in, for instance, a chip set including a processor and a memory as shown in FIG. **14**. As such, the computer vision system **103** and/or the horizon estimation platform **105** can provide means for accomplishing various parts of the process **1000**. In addition or alternatively, the services platform **107** and/or services **109** may perform any combination of the steps of the process **1000** in combination with the computer vision system **103**/the horizon estimation platform **105** or as standalone components. Although the process **1000** is illustrated and described as a sequence of steps, it is contemplated that

various embodiments of the process 1000 may be performed in any order or combination and need not include all of the illustrated steps.

**[0108]** In step 1001, the computer vision system 103 segments an input image into a plurality of grid cells. As previously discussed, the input image can be part of an image capture stream (e.g., from an onboard camera of a vehicle 101) to support visual odometry to more accurately localize the vehicle 101 (e.g., localized to within 10 cm accuracy). In one embodiment, the grid is comprised of regular shapes (e.g., square, rectangle, etc.), but it is contemplated that grid can also vary in size and/or shape from cell to cell. For example, in areas where higher resolution is needed (e.g., in the center of an image), smaller cells can be used to provide greater resolution. Similarly, larger cells can be used for the periphery of an image, where greater resolution may not be needed.

**[0109]** In one embodiment, the resolution or size of the grid can vary with available processing power and/or desired accuracy/preciseness of the resulting lane models. As previously discussed, in one embodiment, the grid resolution is at a relatively low level (e.g., 50×38).

**[0110]** In step 1003, the computer vision system 103 processes a portion of the input image in each of the plurality of grid cells to detect one or more lane lines in said each grid cell. In one embodiment, the computer vision system 103 assigns the processing of said each grid cell or the generating of the parametric representation of said each grid cell to a different processing node of the computer vision system. For example, in a neural network, the portion of the image data falling within a grid cell represent the receptive field for a given collection of neurons. These neurons can then use machine learning to automatically detect lane lines within the image and compute the confidence levels (e.g., prediction confidence) for the detected features. In one embodiment, the computer vision system 103 comprises a convolutional neural network, and the generating of the parametric representation is completed in one forward pass of the convolutional neural network.

**[0111]** In step 1005, the computer vision system 103 determines one or more intercepts of the one or more lane lines with one or more edges of said each grid cell, and one or more slopes of the one or more lane lines at the one or more intercepts for said each grid cell in which the one or more lane lines are detected. In one embodiment, the center line of the detected lanes can be used to determine the edge intercepts and slopes for each grid cell. It is contemplated that the computer vision system 103 can use any coordinate system, measurement unit, and/or scale to define the intercepts and slopes of the line at the intercepts.

**[0112]** In step 1007, the computer vision system 103 generates a parametric representation of the one or more lane lines for said each grid cell, wherein the parametric representation encodes the one or more intercepts and the one or more slopes into a data structure for said each grid cell. In one embodiment, the parametric representation of the data structure includes: (1) an indicator value parameter for each of the one or more edges to indicate which of the one or more edges of said each grid cell the one or more intercepts cross, (2) a slope parameter to indicate a slope of the one or more lane lines at the one or more intercepts, and/or (3) an intercept parameter to indicate a position along the one or more edges at which the one or more intercept occurs. In one embodiment, the indicator value parameter

represents a probability that the one or more lanes is predicted to cross at the one or more intercepts when the computer vision system is operating in a prediction mode. One example of this data structure or parametric representation is discussed with respect to FIG. 6 above.

**[0113]** In one embodiment, the computer vision system 103 can optionally determine that there are a plurality of the one or more lane lines detected for said each grid cell. The computer vision system 103 then generates the parametric representation for each of the plurality of the one or more lane lines, and outputs the parametric representation for said each of the plurality of the one or more lane lines as a different set of output channels. In one embodiment, the computer vision system 103 can be configured with a maximum number of lane lines that it is to detect in a given grid cell. This maximum number can then be used to determine the number of output channels to allocate to the parametric representation. For example, in the example discussed above, a parametric representation of a line can have 12 parameters in one set to describe a single lane line. Accordingly, the number of output channels to allocate be the maximum number lane lines to detect multiplied by the number of parameters in a set.

**[0114]** In one embodiment, the parametric representation can be extended with additional classes and/or attributes to describe a detected lane line. For example, the parametric representation further includes an attribute parameter indicating a lane line type. This lane type can include a description class or attribute of the lane line such as whether the lane line is a solid lane line type or a dashed lane line type. In one embodiment, the lane type can be determined directly from the input image. For example, the computer vision system 103 can identify whether a lane consists of a solid line or a dashed line based on recognized visual features in the input image. In addition or alternatively, the lane type can be inferred or determined from the geographic database 111. For example, the computer vision system 103 can query the database to identify the lane type based on the current coordinates of the vehicle 101 (e.g., as determined from GPS, visual odometry, and/or any other available localization technique).

**[0115]** In step 1009, the computer vision system 103 provides an output parametric representation for the input image, wherein the output parametric representation aggregates the parametric representations of said each grid cell. In other words, the computer vision system 103 can aggregate the parametric representation of lane lines in each grid cell into an output parametric representation that encompasses the entire input image. This collection of the basic representation units at the grid cells can then represent the overall geometry of lane lines or lane models.

**[0116]** FIG. 11 is a flowchart of a process for grouping grid cells based on their respective parametric representations of lane lines, according to one embodiment. In one embodiment, the computer vision system 103 and/or the horizon estimation platform 105 may perform one or more portions of the process 1100 and may be implemented in, for instance, a chip set including a processor and a memory as shown in FIG. 14. As such, the computer vision system 103 and/or the horizon estimation platform 105 can provide means for accomplishing various parts of the process 1100. In addition or alternatively, the services platform 107 and/or services 109 may perform any combination of the steps of the process 1100 in combination with the computer vision

system **103**/the horizon estimation platform **105** or as stand-alone components. Although the process **1100** is illustrated and described as a sequence of steps, it is contemplated that various embodiments of the process **1100** may be performed in any order or combination and need not include all of the illustrated steps.

**[0117]** In one embodiment, the process **1100** is performed following the creation of the parametric representations of at least some or all of the grid cells as described with respect to the process **1000** of FIG. **10**.

**[0118]** In step **1101**, the computer vision system **103** determines shared borders of grid cells. By way of example, a shared border of occurs at an edge of a first grid cell that is immediately adjacent or overlaps a corresponding edge of an adjacent second grid cell. Accordingly, in one embodiment, at each shared border of said each grid cell in the output parametric representation, the computer vision system **103** can look for cells with two edges with large indicator values (step **1103**) and join the intercepts or simply group the cells (step **1107**). For example, a “large indicator” value refers to a predicted or detected lane line for which the computer vision system **103** has predicted a probability of crossing the edge that is above a threshold probability. In other words, when two cells share a common border and each of the two cells has an intercept on the edge at the common border, then the two intercepts or cells can be joined into a common lane line if the indicator value is above the probability threshold.

**[0119]** In addition or alternatively, the computer vision system **103** combines two of said each grid cells (step **1107**) when the one or more intercepts, the one or more slopes, a confidence value associated with the one or more intercepts, or a combination thereof for said two of said each grid cells are within a tolerance value (step **1105**). In this embodiment, instead of relying on just the indicator value to join or group cells as discussed above, the computer vision system **103** can evaluate whether two adjacent intercepts occur at the same position along the shared border within a tolerance level (e.g., a threshold distance) and/or whether the slopes of the two intercepts also match to within a tolerance level (e.g., a threshold degree of variance). If the intercepts and/or slopes match within the tolerance level, then the intercepts or cells are joined into a group.

**[0120]** In one embodiment, a group of said each grid cells resulting from the combining represents a given lane line. In other words, the computationally cheap act of comparing intercepts and/or slopes at shared cell borders can result in building a lane model that advantageously does not require a fully connected layer of a neural network.

**[0121]** In step **1109**, the computer vision system **103** processes the one or more lane lines detected in said each grid cell in the group to generate a polyline representing of a lane model of the given lane.

**[0122]** FIG. **12** is a flowchart for decoding parametric representations of lane lines into polylines, according to one embodiment. In one embodiment, the computer vision system **103** and/or the horizon estimation platform **105** may perform one or more portions of the process **1200** and may be implemented in, for instance, a chip set including a processor and a memory as shown in FIG. **14**. As such, the computer vision system **103** and/or the horizon estimation platform **105** can provide means for accomplishing various parts of the process **1200**. In addition or alternatively, the services platform **107** and/or services **109** may perform any

combination of the steps of the process **1200** in combination with the computer vision system **103**/the horizon estimation platform **105** or as standalone components. Although the process **1200** is illustrated and described as a sequence of steps, it is contemplated that various embodiments of the process **1200** may be performed in any order or combination and need not include all of the illustrated steps.

**[0123]** The process **1200** describes an embodiment of a process for decoding grouped parametric representations of lane lines to generated lane models as described in step **1109** of the process **1100** above. Accordingly, in one embodiment, the process **1200** is performed following the process **1000** of FIG. **10** and the process **1100** of FIG. **11**. One advantage of this decoding approach is that once a parametric representation of a lane line is generated for each grid cell, no further image analysis is needed to manipulate or otherwise decode the lane lines. This, in turn, results in a significant reduction of computational resources needed from create the lane models.

**[0124]** In step **1201**, the processing of the one or more lane lines comprises averaging the one or more slopes, the one or more intercepts, or a combination thereof for said each shared border in the group. As previously discussed, the different intercept and slope results generated by each adjacent cell of shared border represents duplicate information of the detected lane line edge crossing. As a results, averaging the duplicate information or values can advantageously improve the accuracy of the predicted lane lines.

**[0125]** In step **1203**, the computer vision system **103** determines a curvature of the one or more lane lines in said each grid cell based on an input slope and an output slope of the one or more slopes at the one or more intercepts. For example, there are typically two intercepts at two different edges or each grid cell in which a lane line is detected. These two intercepts and their respective slopes can be used to parametrically describe the shape of the lane line shape within the cell without actually have to store any data points about the line other than the intercepts and slope. In other words, no data from the interior of the grid cell is needed. Instead, the computer vision system **103** can compute a curve that encompasses the intercepts with the curvature based on the slopes of the intercepts. This curve can be based on computing, for instance, a Hermite polynomial, Bezier curve, and/or the like.

**[0126]** Once the curve is determined, the computer vision system **103** can initiate the process or converting the curve representations into a polyline or other vector-based representation of the lane lines. At step **1205**, for instance, the computer vision system **103** uses the determined curves to calculate an excess number of points along the curve. As described above, the excess or large number points is a number that is greater than needed to as junction points of a polyline or vector-based representation of the lane line.

**[0127]** At step **1207**, the computer vision system **103** can simplify the point representation of the lane lines by removing any points from the excess points that are not needed to delineate the polyline or vector-based representation to a predetermined accuracy and/or precision. This simplification process can be performed using any algorithm for simplifying curves such as the Ramer-Douglas-Peucker algorithm. The computer vision system **103** then uses the simplified point representation to generate the polyline or vector-based representation of the lane lines by connecting line segments between the remaining points with the remain-

ing points acting as junction points. An example of the resulting polyline is illustrated in the example of FIG. 4A above.

[0128] The processes described herein for providing vanishing point/horizon estimation using lane models may be advantageously implemented via software, hardware (e.g., general processor, Digital Signal Processing (DSP) chip, an Application Specific Integrated Circuit (ASIC), Field Programmable Gate Arrays (FPGAs), etc.), firmware or a combination thereof. Such exemplary hardware for performing the described functions is detailed below.

[0129] FIG. 13 illustrates a computer system 1300 upon which an embodiment of the invention may be implemented. Computer system 1300 is programmed (e.g., via computer program code or instructions) to provide vanishing point/horizon estimation using lane models as described herein and includes a communication mechanism such as a bus 1310 for passing information between other internal and external components of the computer system 1300. Information (also called data) is represented as a physical expression of a measurable phenomenon, typically electric voltages, but including, in other embodiments, such phenomena as magnetic, electromagnetic, pressure, chemical, biological, molecular, atomic, sub-atomic and quantum interactions. For example, north and south magnetic fields, or a zero and non-zero electric voltage, represent two states (0, 1) of a binary digit (bit). Other phenomena can represent digits of a higher base. A superposition of multiple simultaneous quantum states before measurement represents a quantum bit (qubit). A sequence of one or more digits constitutes digital data that is used to represent a number or code for a character. In some embodiments, information called analog data is represented by a near continuum of measurable values within a particular range.

[0130] A bus 1310 includes one or more parallel conductors of information so that information is transferred quickly among devices coupled to the bus 1310. One or more processors 1302 for processing information are coupled with the bus 1310.

[0131] A processor 1302 performs a set of operations on information as specified by computer program code related to providing vanishing point/horizon estimation using lane models. The computer program code is a set of instructions or statements providing instructions for the operation of the processor and/or the computer system to perform specified functions. The code, for example, may be written in a computer programming language that is compiled into a native instruction set of the processor. The code may also be written directly using the native instruction set (e.g., machine language). The set of operations include bringing information in from the bus 1310 and placing information on the bus 1310. The set of operations also typically include comparing two or more units of information, shifting positions of units of information, and combining two or more units of information, such as by addition or multiplication or logical operations like OR, exclusive OR (XOR), and AND. Each operation of the set of operations that can be performed by the processor is represented to the processor by information called instructions, such as an operation code of one or more digits. A sequence of operations to be executed by the processor 1302, such as a sequence of operation codes, constitute processor instructions, also called computer system instructions or, simply, computer instructions. Processors may be implemented as mechanical, electrical, mag-

netic, optical, chemical or quantum components, among others, alone or in combination.

[0132] Computer system 1300 also includes a memory 1304 coupled to bus 1310. The memory 1304, such as a random access memory (RAM) or other dynamic storage device, stores information including processor instructions for providing vanishing point/horizon estimation using lane models. Dynamic memory allows information stored therein to be changed by the computer system 1300. RAM allows a unit of information stored at a location called a memory address to be stored and retrieved independently of information at neighboring addresses. The memory 1304 is also used by the processor 1302 to store temporary values during execution of processor instructions. The computer system 1300 also includes a read only memory (ROM) 1306 or other static storage device coupled to the bus 1310 for storing static information, including instructions, that is not changed by the computer system 1300. Some memory is composed of volatile storage that loses the information stored thereon when power is lost. Also coupled to bus 1310 is a non-volatile (persistent) storage device 1308, such as a magnetic disk, optical disk or flash card, for storing information, including instructions, that persists even when the computer system 1300 is turned off or otherwise loses power.

[0133] Information, including instructions for providing vanishing point/horizon estimation using lane models, is provided to the bus 1310 for use by the processor from an external input device 1312, such as a keyboard containing alphanumeric keys operated by a human user, or a sensor. A sensor detects conditions in its vicinity and transforms those detections into physical expression compatible with the measurable phenomenon used to represent information in computer system 1300. Other external devices coupled to bus 1310, used primarily for interacting with humans, include a display device 1314, such as a cathode ray tube (CRT) or a liquid crystal display (LCD), or plasma screen or printer for presenting text or images, and a pointing device 1316, such as a mouse or a trackball or cursor direction keys, or motion sensor, for controlling a position of a small cursor image presented on the display 1314 and issuing commands associated with graphical elements presented on the display 1314. In some embodiments, for example, in embodiments in which the computer system 1300 performs all functions automatically without human input, one or more of external input device 1312, display device 1314 and pointing device 1316 is omitted.

[0134] In the illustrated embodiment, special purpose hardware, such as an application specific integrated circuit (ASIC) 1320, is coupled to bus 1310. The special purpose hardware is configured to perform operations not performed by processor 1302 quickly enough for special purposes. Examples of application specific ICs include graphics accelerator cards for generating images for display 1314, cryptographic boards for encrypting and decrypting messages sent over a network, speech recognition, and interfaces to special external devices, such as robotic arms and medical scanning equipment that repeatedly perform some complex sequence of operations that are more efficiently implemented in hardware.

[0135] Computer system 1300 also includes one or more instances of a communications interface 1370 coupled to bus 1310. Communication interface 1370 provides a one-way or two-way communication coupling to a variety of external

devices that operate with their own processors, such as printers, scanners and external disks. In general the coupling is with a network link **1378** that is connected to a local network **1380** to which a variety of external devices with their own processors are connected. For example, communication interface **1370** may be a parallel port or a serial port or a universal serial bus (USB) port on a personal computer. In some embodiments, communications interface **1370** is an integrated services digital network (ISDN) card or a digital subscriber line (DSL) card or a telephone modem that provides an information communication connection to a corresponding type of telephone line. In some embodiments, a communication interface **1370** is a cable modem that converts signals on bus **1310** into signals for a communication connection over a coaxial cable or into optical signals for a communication connection over a fiber optic cable. As another example, communications interface **1370** may be a local area network (LAN) card to provide a data communication connection to a compatible LAN, such as Ethernet. Wireless links may also be implemented. For wireless links, the communications interface **1370** sends or receives or both sends and receives electrical, acoustic or electromagnetic signals, including infrared and optical signals, that carry information streams, such as digital data. For example, in wireless handheld devices, such as mobile telephones like cell phones, the communications interface **1370** includes a radio band electromagnetic transmitter and receiver called a radio transceiver. In certain embodiments, the communications interface **1370** enables connection to the communication network **113** for providing vanishing point/horizon estimation using lane models.

**[0136]** The term computer-readable medium is used herein to refer to any medium that participates in providing information to processor **1302**, including instructions for execution. Such a medium may take many forms, including, but not limited to, non-volatile media, volatile media and transmission media. Non-volatile media include, for example, optical or magnetic disks, such as storage device **1308**. Volatile media include, for example, dynamic memory **1304**. Transmission media include, for example, coaxial cables, copper wire, fiber optic cables, and carrier waves that travel through space without wires or cables, such as acoustic waves and electromagnetic waves, including radio, optical and infrared waves. Signals include man-made transient variations in amplitude, frequency, phase, polarization or other physical properties transmitted through the transmission media. Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, any other magnetic medium, a CD-ROM, CDRW, DVD, any other optical medium, punch cards, paper tape, optical mark sheets, any other physical medium with patterns of holes or other optically recognizable indicia, a RAM, a PROM, an EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave, or any other medium from which a computer can read.

**[0137]** FIG. **14** illustrates a chip set **1400** upon which an embodiment of the invention may be implemented. Chip set **1400** is programmed to provide vanishing point/horizon estimation using lane models as described herein and includes, for instance, the processor and memory components described with respect to FIG. **13** incorporated in one or more physical packages (e.g., chips). By way of example, a physical package includes an arrangement of one or more materials, components, and/or wires on a structural assem-

bly (e.g., a baseboard) to provide one or more characteristics such as physical strength, conservation of size, and/or limitation of electrical interaction. It is contemplated that in certain embodiments the chip set can be implemented in a single chip.

**[0138]** In one embodiment, the chip set **1400** includes a communication mechanism such as a bus **1401** for passing information among the components of the chip set **1400**. A processor **1403** has connectivity to the bus **1401** to execute instructions and process information stored in, for example, a memory **1405**. The processor **1403** may include one or more processing cores with each core configured to perform independently. A multi-core processor enables multiprocessing within a single physical package. Examples of a multi-core processor include two, four, eight, or greater numbers of processing cores. Alternatively or in addition, the processor **1403** may include one or more microprocessors configured in tandem via the bus **1401** to enable independent execution of instructions, pipelining, and multithreading. The processor **1403** may also be accompanied with one or more specialized components to perform certain processing functions and tasks such as one or more digital signal processors (DSP) **1407**, or one or more application-specific integrated circuits (ASIC) **1409**. A DSP **1407** typically is configured to process real-world signals (e.g., sound) in real time independently of the processor **1403**. Similarly, an ASIC **1409** can be configured to performed specialized functions not easily performed by a general purposed processor. Other specialized components to aid in performing the inventive functions described herein include one or more field programmable gate arrays (FPGA) (not shown), one or more controllers (not shown), or one or more other special-purpose computer chips.

**[0139]** The processor **1403** and accompanying components have connectivity to the memory **1405** via the bus **1401**. The memory **1405** includes both dynamic memory (e.g., RAM, magnetic disk, writable optical disk, etc.) and static memory (e.g., ROM, CD-ROM, etc.) for storing executable instructions that when executed perform the inventive steps described herein to provide vanishing point/horizon estimation using lane models. The memory **1405** also stores the data associated with or generated by the execution of the inventive steps.

**[0140]** FIG. **15** is a diagram of exemplary components of a mobile station (e.g., handset) capable of operating in the system of FIG. **1**, according to one embodiment. Generally, a radio receiver is often defined in terms of front-end and back-end characteristics. The front-end of the receiver encompasses all of the Radio Frequency (RF) circuitry whereas the back-end encompasses all of the base-band processing circuitry. Pertinent internal components of the telephone include a Main Control Unit (MCU) **1503**, a Digital Signal Processor (DSP) **1505**, and a receiver/transmitter unit including a microphone gain control unit and a speaker gain control unit. A main display unit **1507** provides a display to the user in support of various applications and mobile station functions that offer automatic contact matching. An audio function circuitry **1509** includes a microphone **1511** and microphone amplifier that amplifies the speech signal output from the microphone **1511**. The amplified speech signal output from the microphone **1511** is fed to a coder/decoder (CODEC) **1513**.

**[0141]** A radio section **1515** amplifies power and converts frequency in order to communicate with a base station,

which is included in a mobile communication system, via antenna 1517. The power amplifier (PA) 1519 and the transmitter/modulation circuitry are operationally responsive to the MCU 1503, with an output from the PA 1519 coupled to the duplexer 1521 or circulator or antenna switch, as known in the art. The PA 1519 also couples to a battery interface and power control unit 1520.

[0142] In use, a user of mobile station 1501 speaks into the microphone 1511 and his or her voice along with any detected background noise is converted into an analog voltage. The analog voltage is then converted into a digital signal through the Analog to Digital Converter (ADC) 1523. The control unit 1503 routes the digital signal into the DSP 1505 for processing therein, such as speech encoding, channel encoding, encrypting, and interleaving. In one embodiment, the processed voice signals are encoded, by units not separately shown, using a cellular transmission protocol such as global evolution (EDGE), general packet radio service (GPRS), global system for mobile communications (GSM), Internet protocol multimedia subsystem (IMS), universal mobile telecommunications system (UMTS), etc., as well as any other suitable wireless medium, e.g., microwave access (WiMAX), Long Term Evolution (LTE) networks, code division multiple access (CDMA), wireless fidelity (WiFi), satellite, and the like.

[0143] The encoded signals are then routed to an equalizer 1525 for compensation of any frequency-dependent impairments that occur during transmission through the air such as phase and amplitude distortion. After equalizing the bit stream, the modulator 1527 combines the signal with a RF signal generated in the RF interface 1529. The modulator 1527 generates a sine wave by way of frequency or phase modulation. In order to prepare the signal for transmission, an up-converter 1531 combines the sine wave output from the modulator 1527 with another sine wave generated by a synthesizer 1533 to achieve the desired frequency of transmission. The signal is then sent through a PA 1519 to increase the signal to an appropriate power level. In practical systems, the PA 1519 acts as a variable gain amplifier whose gain is controlled by the DSP 1505 from information received from a network base station. The signal is then filtered within the duplexer 1521 and optionally sent to an antenna coupler 1535 to match impedances to provide maximum power transfer. Finally, the signal is transmitted via antenna 1517 to a local base station. An automatic gain control (AGC) can be supplied to control the gain of the final stages of the receiver. The signals may be forwarded from there to a remote telephone which may be another cellular telephone, other mobile phone or a land-line connected to a Public Switched Telephone Network (PSTN), or other telephony networks.

[0144] Voice signals transmitted to the mobile station 1501 are received via antenna 1517 and immediately amplified by a low noise amplifier (LNA) 1537. A down-converter 1539 lowers the carrier frequency while the demodulator 1541 strips away the RF leaving only a digital bit stream. The signal then goes through the equalizer 1525 and is processed by the DSP 1505. A Digital to Analog Converter (DAC) 1543 converts the signal and the resulting output is transmitted to the user through the speaker 1545, all under control of a Main Control Unit (MCU) 1503—which can be implemented as a Central Processing Unit (CPU) (not shown).

[0145] The MCU 1503 receives various signals including input signals from the keyboard 1547. The keyboard 1547 and/or the MCU 1503 in combination with other user input components (e.g., the microphone 1511) comprise a user interface circuitry for managing user input. The MCU 1503 runs a user interface software to facilitate user control of at least some functions of the mobile station 1501 to provide vanishing point/horizon estimation using lane models. The MCU 1503 also delivers a display command and a switch command to the display 1507 and to the speech output switching controller, respectively. Further, the MCU 1503 exchanges information with the DSP 1505 and can access an optionally incorporated SIM card 1549 and a memory 1551. In addition, the MCU 1503 executes various control functions required of the station. The DSP 1505 may, depending upon the implementation, perform any of a variety of conventional digital processing functions on the voice signals. Additionally, DSP 1505 determines the background noise level of the local environment from the signals detected by microphone 1511 and sets the gain of microphone 1511 to a level selected to compensate for the natural tendency of the user of the mobile station 1501.

[0146] The CODEC 1513 includes the ADC 1523 and DAC 1543. The memory 1551 stores various data including call incoming tone data and is capable of storing other data including music data received via, e.g., the global Internet. The software module could reside in RAM memory, flash memory, registers, or any other form of writable computer-readable storage medium known in the art including non-transitory computer-readable storage medium. For example, the memory device 1551 may be, but not limited to, a single memory, CD, DVD, ROM, RAM, EEPROM, optical storage, or any other non-volatile or non-transitory storage medium capable of storing digital data.

[0147] An optionally incorporated SIM card 1549 carries, for instance, important information, such as the cellular phone number, the carrier supplying service, subscription details, and security information. The SIM card 1549 serves primarily to identify the mobile station 1501 on a radio network. The card 1549 also contains a memory for storing a personal telephone number registry, text messages, and user specific mobile station settings.

[0148] While the invention has been described in connection with a number of embodiments and implementations, the invention is not so limited but covers various obvious modifications and equivalent arrangements, which fall within the purview of the appended claims. Although features of the invention are expressed in certain combinations among the claims, it is contemplated that these features can be arranged in any combination and order.

What is claimed is:

1. A computer-implemented method for determining a horizon in an image depicting one or more lanes of a roadway comprising:

processing the image to construct one or more lane models of the one or more road lanes depicted in the image;

extending the one or more road lanes through the image using the one or more lane models;

determining a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line; and

- designating the horizontal line as the horizon of the image.
2. The method of claim 1, further comprising: initiating a line sweep of the image from a top to a bottom of the image to determine the horizontal line.
  3. The method of claim 1, wherein the one or more road lanes are extended towards a top of the image.
  4. The method of claim 1, further comprising: detecting a slope of the roadway; adjusting the horizon based on the slope.
  5. The method of claim 4, further comprising: determining the slope from one or more sensors of a vehicle, wherein the vehicle is equipped with a camera sensor to capture the image as it travels on the roadway.
  6. The method of claim 1, further comprising: using the horizon to determine a consistency of the one or more lane models detected in one or more subsequent images of the one or more lanes the roadway.
  7. The method of claim 1, further comprising: estimating a transformation of a perspective representation to an orthographic projection of the one or more road lanes based on the horizon and the one or more lane models.
  8. The method of claim 1, further comprising: labeling or annotating one or more map features depicted in the image based on the horizon, wherein localization of a vehicle is further based on the labeling or the annotating.
  9. The method of claim 1, further comprising: selecting a portion of the image for performing a sky detection based on the horizon.
  10. An apparatus for determining a horizon in an image depicting one or more lanes of a roadway comprising: at least one processor; and at least one memory including computer program code for one or more programs, the at least one memory and the computer program code configured to, with the at least one processor, cause the apparatus to perform at least the following, process the image to construct one or more lane models of the one or more road lanes depicted in the image; extend the one or more road lanes through the image using the one or more lane models; determine a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line; and designate the horizontal line as the horizon of the image.
  11. The apparatus of claim 10, wherein the apparatus is further caused to: initiate a line sweep of the image from a top to a bottom of the image to determine the horizontal line.
  12. The apparatus of claim 10, wherein the one or more road lanes are extended towards a top of the image.
  13. The apparatus of claim 10, wherein the apparatus is further caused to: detect a slope of the roadway; adjust the horizon based on the slope.
  14. The apparatus of claim 13, wherein the apparatus is further caused to: determine the slope from one or more sensors of a vehicle, wherein the vehicle is equipped with a camera sensor to capture the image as it travels on the roadway.
  15. The apparatus of claim 10, wherein the apparatus is further caused to perform at least one of the following: use the horizon to determine a consistency of the one or more lane models detected in one or more subsequent images of the one or more lanes the roadway; estimate a transformation of a perspective representation to an orthographic projection of the one or more road lanes based on the horizon and the one or more lane models; label or annotate one or more map features depicted in the image based on the horizon, wherein localization of a vehicle is further based on the labeling or the annotating; and select a portion of the image for performing a sky detection based on the horizon.
  16. A non-transitory computer-readable storage medium for determining a vanishing point in an image depicting one or more lanes of a roadway, carrying one or more sequences of one or more instructions which, when executed by one or more processors, cause an apparatus to perform: processing the image to construct one or more lane models of the one or more road lanes depicted in the image; extending the one or more road lanes through the image using the one or more lane models; determining a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line; and designating the horizontal line as the vanishing point of the image.
  17. The non-transitory computer-readable storage medium of claim 16, wherein the apparatus is further caused to perform: initiating a line sweep of the image from a top to a bottom of the image to determine the horizontal line.
  18. The non-transitory computer-readable storage medium of claim 16, wherein the apparatus is further caused to perform: detecting a slope of the roadway; adjusting the vanishing point based on the slope.
  19. The non-transitory computer-readable storage medium of claim 18, wherein the apparatus is further caused to perform: determining the slope from one or more sensors of a vehicle, wherein the vehicle is equipped with a camera sensor to capture the image as it travels on the roadway.
  20. The method of claim 10, wherein the apparatus is further caused to perform at least one of the following: using the vanishing point to determine a consistency of the one or more lane models detected in one or more subsequent images of the one or more lanes the roadway; estimating a transformation of a perspective representation to an orthographic projection of the one or more road lanes based on the vanishing point and the one or more lane models;

labeling or annotating one or more map features depicted in the image based on the vanishing point, wherein localization of a vehicle is further based on the labeling or the annotating; and selecting a portion of the image for performing a sky detection based on the vanishing point.

\* \* \* \* \*