



(51) International Patent Classification:

G06K 9/00 (2006.01)

G06K 9/62 (2006.01)

G06K 9/66 (2006.01)

(21) International Application Number:

PCT/US2019/022592

(22) International Filing Date:

15 March 2019 (15.03.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/643,665 15 March 2018 (15.03.2018) US

16/355,328 15 March 2019 (15.03.2019) US

(71) Applicant: NVIDIA CORPORATION [US/US]; 2530 Zanker Road, San Jose, California 95131 (US).

(72) Inventors: RANKAWAT, Mansi; c/o Nvidia Corporation, 2530 Zanker Road, San Jose, California 95131 (US). YAO,

Jian; c/o Nvidia Corporation, 2530 Zanker Road, San Jose, California 95131 (US). ZHANG, Dong; c/o Nvidia Corporation, 2530 Zanker Road, San Jose, California 95131 (US). CHEN, Chia-Chih; c/o Nvidia Corporation, 2530 Zanker Road, San Jose, California 95131 (US).

(74) Agent: BACON, Kirk D. et al.; Shook, Hardy & Bacon, LLP, 2555 Grand Blvd., Kansas City, Missouri 64108 (US).

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

(54) Title: DETERMINING DRIVABLE FREE-SPACE FOR AUTONOMOUS VEHICLES




FIGURE 1A

(57) Abstract: In various examples, sensor data may be received that represents a field of view of a sensor of a vehicle located in a physical environment. The sensor data may be applied to a machine learning model that computes both a set of boundary points that correspond to a boundary dividing drivable free-space from non-drivable space in the physical environment and class labels for boundary points of the set of boundary points that correspond to the boundary. Locations within the physical environment may be determined from the set of boundary points represented by the sensor data, and the vehicle may be controlled through the physical environment within the drivable free-space using the locations and the class labels.

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

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

**Published:**

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

## DETERMINING DRIVABLE FREE-SPACE FOR AUTONOMOUS VEHICLES

### BACKGROUND OF THE INVENTION

An autonomous driving system should control an autonomous vehicle without human supervision while achieving an acceptable level of safety. This may require the autonomous driving system to be capable of achieving at least the functional performance of an attentive human driver, who draws upon a perception and action system that has an incredible ability to identify and react to moving and static obstacles in a complex environment. In order to accomplish this, areas of the environment that are obstacle-free (e.g., drivable free-space) may be determined, as this information may be useful to the autonomous driving system and/or advanced driver assistance systems (ADAS) when planning maneuvers and/or navigation decisions.

Some conventional approaches to determining drivable free-space have used vision-based techniques using deep artificial neural networks (DNN). For example, these conventional approaches have used a DNN, such as a convolutional neural network (CNN), to perform semantic segmentation (e.g., pixel-wise classification of an image). However, semantic segmentation may be computationally expensive because a classification is assigned to each pixel of an image and, as a result, may require extensive post-processing on outputs of the DNN to make the outputs useable by the autonomous driving system. Thus, a drawback of many conventional approaches that use semantic segmentation for determining drivable free-space is their inability to run in real-time. In some conventional approaches – such as where adjustments are made to reduce the computational expense to allow real-time operation – semantic segmentation comes at the expense of determining the drivable free-space below a level of accuracy required to maintain an acceptable level of safety for autonomous driving.

In addition, in other conventional approaches, a CNN may be implemented that may perform column-wise regression. However, these conventional approaches may use fully connected layers which, similar to the semantic segmentation tasks described above, also may consume an excessive amount of computing resources, thereby reducing the ability of the CNNs to run in real-time. In addition, even where CNNs are used to perform column-wise regression, the type or class of boundary or barrier regressed upon in each column is not identified. As a result, the output of the CNNs may not be informative enough to an

- 2 -

autonomous driving system to enable safe operation of an autonomous vehicle. For example, without context provided by the CNN of a class of boundary – e.g., a dynamic boundary, such as a human, as opposed to a static boundary, such as a curb – the autonomous driving system may not be able to accurately predict drivable free-space in a way that results in safe control of the autonomous vehicle. In such an example, when determining where to navigate, the autonomous driving system may not take into account the dynamic nature of the boundary (e.g., a dynamic boundary class may decrease the drivable free-space upon movement). As a result, the autonomous vehicle may use the determined drivable free-space for navigating the environment even where the determined drivable free-space does not correspond to the actual drivable free-space (e.g., as a result of movement of one or more dynamic boundary classes).

## SUMMARY OF THE INVENTION

Embodiments of the present disclosure relate to determining drivable free-space for autonomous vehicles. More specifically, systems and methods are disclosed for identifying one or more boundaries separating drivable free-space (e.g., obstacle free-space) from non-drivable space (e.g., space with or beyond one or more obstacles) in a physical environment for use by an autonomous vehicle in navigating the physical environment.

In contrast to conventional systems, such as those described above, systems of the present disclosure may use an efficient and accurate machine learning model – such as a convolutional neural network (CNN) – to regress on one or more boundaries separating drivable free-space from non-drivable space in a physical environment. For example, the CNN of present systems may be a fully convolutional network, meaning the CNN may not include any fully-connected layers, thereby increasing the efficiency of the systems while reducing the drain on computing resources. In addition, by regressing on the boundary(ies) (e.g., column by column), the current systems may not require separately classifying each pixel of an image – as required by conventional segmentation approaches – thereby reducing the requirement of performing extensive post-processing on the output of the CNN.

In further contrast to conventional approaches, the CNN of the present systems may predict labels for each of the boundary classes corresponding to the boundary(ies) identified in the image. As a result, the present systems may use this contextual information (e.g., dynamic boundary, static boundary, vehicle, pedestrian, curb, barrier, etc.) to navigate an autonomous vehicle through the environment safely, taking into consideration the different

boundary classes delineating the drivable free-space – such as whether the boundary classes may move, how they may move, where they move to, and/or the like.

Ultimately, the present systems may implement a CNN for detecting drivable free-space that – compared to conventional approaches – is computationally less expensive, more contextually informative, efficient enough to run in real-time (e.g., at 30 frames per second or greater), and accurate enough for use in navigating an autonomous vehicle through a real-world physical environment safely.

### BRIEF DESCRIPTION OF THE DRAWING

The present systems and methods for determining drivable free-space for autonomous vehicles is described in detail below with reference to the attached drawing figures, wherein:

FIG. 1A is an illustration of a data flow diagram for boundary identification, in accordance with some embodiments of the present disclosure;

FIG. 1B is an illustration of an example machine learning model for boundary identification, in accordance with some embodiments of the present disclosure;

FIG. 1C is another illustration of an example machine learning model for boundary identification, in accordance with some embodiments of the present disclosure;

FIG. 1D is an illustration of an example output of a machine learning model for boundary identification, in accordance with some embodiments of the present disclosure;

FIG. 2 is an example flow diagram for a method of determining a boundary, in accordance with some embodiments of the present disclosure;

FIG. 3A is an example training image for a machine learning model, in accordance with some embodiments of the present disclosure;

FIG. 3B is another example training image for a machine learning model, in accordance with some embodiments of the present disclosure;

FIG. 4 is an example flow diagram for a method of training a machine learning model, in accordance with some embodiments of the present disclosure;

FIG. 5 is another example training image for a machine learning model, in accordance with some embodiments of the present disclosure;

FIG. 6 is an example flow diagram for a method of detecting multiple boundaries, in accordance with some embodiments of the present disclosure;

FIG. 7A is an illustration of an example autonomous vehicle, in accordance with some embodiments of the present disclosure;

FIG. 7B is an example of camera locations and fields of view for the example autonomous vehicle of FIG. 7A, in accordance with some embodiments of the present disclosure;

FIG. 7C is a block diagram of an example system architecture for the example autonomous vehicle of FIG. 7A, in accordance with some embodiments of the present disclosure;

FIG. 7D is a system diagram for communication between cloud-based server(s) and the example autonomous vehicle of FIG. 7A, in accordance with some embodiments of the present disclosure; and

FIG. 8 is a block diagram of an example computing device suitable for use in implementing some embodiments of the present disclosure.

#### DETAILED DESCRIPTION OF THE INVENTION

Systems and methods are disclosed related to determining drivable free-space for autonomous vehicles. The present disclosure may be described with respect to an example autonomous vehicle 700 (alternatively referred to herein as “vehicle 700” or “autonomous vehicle 700”), an example of which is described in more detail herein with respect to FIGs. 7A-7D. However, the disclosure is not limited to autonomous vehicles, and may be used in Advanced Driver Assistance Systems (ADAS), robotics, virtual reality (e.g., to determine free-space for movement of a player), and/or in other technology areas. As such, the description herein with respect to the vehicle 700 is for example purposes only, and is not intended to limit the disclosure to any one technology area.

##### **Boundary Point and Class Label Detection System**

As described herein, some conventional approaches to determining drivable free-space may rely on deep neural networks (DNN) for generating pixel-by-pixel segmentation information. However, segmentation approaches for determining drivable free-space are computationally expensive and require extensive post-processing in order for an autonomous driving system to make sense of and use the DNN output.

In conventional systems, column-wise regression may be used to determine obstacle positions for identifying drivable free-space. However, these conventional systems

do not account for one or more boundary types or classes that make up or delineate the boundary, thereby minimizing the usefulness of the boundary location output by the DNN. For example, without information as to the boundary class (e.g., static, dynamic, vehicle, barrier, pedestrian, etc.), the location of the boundary may not be useful enough information to safely determine or react to drivable free-space (e.g., pedestrians, vehicles, and other dynamic objects may change location, while static objects may not). In addition, these conventional approaches use DNNs that require significant computing resources to implement and, as a result, may not operate effectively in real-time implementations. For example, conventional DNNs may use multiple fully-connected layers – each of which causes a drain on computing resources – in order to accurately regress on a boundary. As a result, even where a boundary is identified, the process may be unable to perform in real-time and – because the boundary classes may not be identified – the output may not be usable by the system in a way that allows the system to perform as safely as desired (e.g., because the system may not take into account the variable of dynamic boundary classes, such as pedestrians, vehicle, animals, etc.).

In contrast to these conventional systems, systems of the present disclosure may use a DNN – such as a convolutional neural network (CNN) – to regress on one or more boundary(ies) determined to divide drivable free-space from non-drivable space. In addition, the DNN may predict one or more boundary classes associated with the boundary(ies). By using regression rather than segmentation or classification in addition to performing the regression using a comparatively smaller footprint, less processing intensive DNN – such as by not including any fully connected layers, by including less fully connected layers than conventional approaches, by using a fully convolutional network, and/or the like – the DNN of present systems may be able to operate in real-time (e.g., at 30 frames per second (fps) or greater) while also outputting accurate and useable information. For example, the output of the DNN may be used by one or more layers of an autonomous driving software stack, such as a perception layer, a world model management layer, a planning layer, a control layer, an actuation layer, and/or another layer. As a result of the output including not only a location of the boundary, but also the boundary classes, the layer(s) of the autonomous driving software stack may be able to more effectively use the location by accounting for changes in the boundary as a result of one or more variables associated with the boundary class(es) – such as whether the boundary class(es) are static or dynamic, a person or a vehicle, a curb or a cement divider, etc.

In some examples, the output of the DNN may undergo one or more smoothing processes – such as spatial smoothing and/or temporal smoothing. As described herein, because regression may be used instead of segmentation, less complex and robust post-processing techniques may be used while achieving more accurate results. For example, spatial smoothing and/or temporal smoothing may require relatively less computational resources than conventional post-processing techniques for segmentation based approaches, and implementing spatial smoothing and/or temporal smoothing may be comparatively less complex.

Temporal smoothing may be applied to both boundary point locations and boundary class labels output by the DNN. Temporal smoothing may be used to increase the consistency of the result over time. Predicted boundary point locations and/or boundary class labels from one or more previous frames (e.g., image frames represented by sensor data) may be used to compute updated, or smooth, boundary point locations and/or boundary class labels of a current frame. As a result, boundary point locations and/or boundary class labels of a current frame may benefit from prior predictions of boundary point locations and/or boundary class labels, thereby resulting in a smoother, less noisy, output for use by the layer(s) of the autonomous driving software stack.

Spatial smoothing with a Gaussian filter may be applied to predicted boundary point locations. As a result, abrupt changes in the boundary point locations may be eliminated, or smoothed. For example, for a predicted boundary point location at a column of a current frame, values from one or more adjacent columns (e.g., as determined by the Gaussian filter radius) may be used to update, or smooth, the predicted boundary point location of the column. This process may be used for any number of the columns of the current frame, and may result in a smoother, less noisy, output for use by the layer(s) of the autonomous driving software stack.

The predicted boundary point locations may represent pixel locations within an image represented by the sensor data. The pixel locations may be two-dimensional (2D) coordinates in the image (e.g., a column and a row). The 2D coordinates may be converted to three-dimensional (3D) or 2D real-world coordinates corresponding to locations in the physical environment of the vehicle (e.g., global positioning system (GPS) coordinates or other global navigation satellite system (GNSS) coordinates). In order to accurately determine the relationship between the 2D coordinates and the real-world coordinates, 3D to 2D projection may be used. For example, a camera or other sensor(s) may be calibrated

- 7 -

using one or more intrinsic (e.g., focal length,  $f$ , optical center ( $u_o$ ,  $v_o$ ), pixel aspect ratio,  $\alpha$ , skew,  $s$ , etc.) and/or extrinsic (e.g., 3D rotation,  $R$ , translation,  $t$ , etc.) camera parameters. One or more constraints may also be imposed, such as requiring that the 3D point always lies on the ground plane of the driving surface (e.g., because the boundary delineating drivable free-space may be part of, or may extend along, the driving surface). In some examples, one or more of the parameters of the camera may be dynamic (e.g., due to vibration, movement, orientation, etc.), and the 3D to 2D projection may be dynamically updated as a result. In some examples, such as where two or more cameras are used, stereo vision techniques may be used to determine a correlation between 2D points and 3D real-world locations.

In any example, the real-world coordinates may then be mapped to the 2D coordinates of the pixels in the image, such that when the boundary points are determined, the real-world coordinates are known and may be used by the autonomous driving software stack (or more generally, by the autonomous vehicle). More specifically, a distance from a camera center to the boundary in the real-world environment may be determined, and the autonomous vehicle may use the distance to each of the boundary points as a drivable free-space in which to operate.

When training the DNN, ground truth data may be generated for use in the training. For example, annotations may be generated (e.g., by a human, by a machine, or a combination thereof). The annotations may include a label for a pixel for each column within a training image. The pixel may be the first pixel in that particular column (beginning from the bottom of the image and moving to the top of the image -- e.g., from closest to furthest from the vehicle in the real-world environment), that corresponds to a boundary that divides drivable free-space from non-drivable space. The annotations may also include indications of classes, which may be indicated by differing colors, shapes, and/or other attributes of the labels. As a result, the DNN may learn to identify not only boundary point locations, but also associated boundary class labels for each of the boundary point locations. In some examples, such as where two or more boundaries are annotated, there may be an annotation on a number of pixels in each column of the image that corresponds to the number of boundaries the DNN is being trained to detect. In addition, the different boundaries may include an additional label, such as drivable (or traversable) or non-drivable (or non-traversable). For example, a curb in front of a driveway or sidewalk may be a traversable boundary while a curb immediately before a building may be a non-traversable boundary. As such, the DNN may learn to differentiate drivable from non-drivable boundaries such that, in emergency

situations or the like, the autonomous vehicle may determine to traverse the drivable boundary to avoid collision with an object or boundary. The DNN is not limited to computing traversable and non-traversable boundaries, and may additionally, or alternatively, compute any number of boundaries having any number of different class assignments.

5                   During training, one or more loss functions may be used to compare the outputs or predictions of the DNN with the ground truth data. For example, a first (boundary) loss function may be used for the boundary point locations output by the DNN and a second (label) loss function may be used for the class labels output by the DNN. In some examples, the boundary loss function may include an L1 loss function and the label loss function may  
10                   include a cross-entropy loss function. A weighted combination of the boundary loss function and the label loss function may be used for a final loss computation. In some examples, a weight ratio may be used between the boundary loss function and the weight loss function (e.g., 1:10, where 1 is for boundary and 10 is for label). An auto-weight scheme may be used that adjusts the cost weights adaptively or automatically when given target ratios for the  
15                   weighted or final losses. The auto-weight scheme may be updated persistently, periodically (e.g., once per training epoch), and/or at another interval, to set the weights.

                  During training, one or more metrics may be used to evaluate the performance of a present system. For example, relative gap, precision, label accuracy, smoothness, weighted free-space precision (WFP), and/or weighted free-space recall (WFR) may be used.  
20                   Relative gap may measure the average of the absolute deviation between a ground truth boundary curve and a predicted boundary curve. Precision may measure the amount of overshoot of the predicted boundary curve from the ground truth boundary curve (e.g., the amount of non-drivable space determined to be drivable free-space by the DNN). Label accuracy may measure the accuracy of the predicted boundary class labels in comparison to  
25                   the ground truth boundary class labels. Smoothness may measure the smoothness of the predicted curve (e.g., the mean of the difference between consecutive boundary points of the boundary curve). WFP may include a precision calculation using only overshoot, such that if there is no overshoot, the precision may be 1.0 (e.g., a near perfect match between ground truth and prediction). WFR may include a precision calculation using only undershoot, such  
30                   that if there is no undershoot, the precision may be 1.0 (e.g., a near perfect match between ground truth and prediction). For WFP and WFR, in some examples, the weight for different boundary class labels may be different, such that some boundary classes or types have a

higher associated weight. For example, pedestrians may have a higher associated weight than a curb because accuracy with respect to a human is more pertinent than with respect to a curb.

Now referring to FIG. 1A, FIG. 1A is an illustration of a data flow diagram for a process 100 of boundary identification, in accordance with some embodiments of the present disclosure. The process 100 for boundary identification may include generating and/or receiving sensor data 102 from one or more sensors of the autonomous vehicle 700. The sensor data 102 may include sensor data from any of the sensors of the vehicle 700 (and/or other vehicles or objects, such as robotic devices, VR systems, AR systems, etc., in some examples). With reference to FIGs. 7A-7C, the sensor data 102 may include the data generated by, for example and without limitation, global navigation satellite systems (GNSS) sensor(s) 758 (e.g., Global Positioning System sensor(s)), RADAR sensor(s) 760, ultrasonic sensor(s) 762, LIDAR sensor(s) 764, inertial measurement unit (IMU) sensor(s) 766 (e.g., accelerometer(s), gyroscope(s), magnetic compass(es), magnetometer(s), etc.), microphone(s) 796, stereo camera(s) 768, wide-view camera(s) 770 (e.g., fisheye cameras), infrared camera(s) 772, surround camera(s) 774 (e.g., 360 degree cameras), long-range and/or mid-range camera(s) 798, speed sensor(s) 744 (e.g., for measuring the speed of the vehicle 700), vibration sensor(s) 742, steering sensor(s) 740, brake sensor(s) (e.g., as part of the brake sensor system 746), and/or other sensor types.

In some examples, the sensor data 102 may include the sensor data generated by one or more forward-facing cameras (e.g., a center or near-center mounted camera(s)) when the vehicle 700 is moving forward, such as a wide-view camera 770, a surround camera 774, a stereo camera 768, and/or a long-range or mid-range camera 798. When the vehicle 700 is moving in reverse, one or more rear-facing cameras may be used. In some examples, one or more side-facing cameras and/or one or more parking cameras may be used to determine drivable free-space to the side of the vehicle 700 and/or immediately adjacent the vehicle 700. More than one camera or other sensor (e.g., LIDAR sensor, RADAR sensor, etc.) may be used to incorporate multiple fields of view (e.g., the fields of view of the long-range cameras 798, the forward-facing stereo camera 768, and/or the forward facing wide-view camera 770 of FIG. 7B). In any example, the sensor data 102 used by machine learning model(s) 104 may be the sensor data 102 determined to be most useful for determining drivable free-space and non-drivable space in a current direction of travel of the vehicle 700.

The sensor data 102 may include image data representing an image(s), image data representing a video (e.g., snapshots of video), and/or sensor data representing fields of

view of sensors (e.g., LIDAR data from LIDAR sensor(s) 764, RADAR data from RADAR sensor(s) 760, etc.). In some examples, the sensor data 102 may be input into the machine learning model(s) 104 and used by the machine learning model(s) 104 to compute boundary points 106 and/or class labels 108. In some other examples, the sensor data 102 may be provided as input to a sensor data pre-processor (not shown) to generate pre-processed sensor data. The pre-processed sensor data may then be input into the machine learning model(s) 104 as input data in addition to, or alternatively from, the sensor data 102.

In examples where the sensor data 102 is image data, many types of images or formats may be used as inputs. For example, compressed images such as in Joint Photographic Experts Group (JPEG) or Luminance/Chrominance (YUV) formats, compressed images as frames stemming from a compressed video format such as H.264/Advanced Video Coding (AVC) or H.265/High Efficiency Video Coding (HEVC), raw images such as originating from Red Clear Blue (RCCB), Red Clear (RCCC) or other type of imaging sensor. It is noted that different formats and/or resolutions could be used training the machine learning model(s) 104 than for inferencing (e.g., during deployment of the machine learning model(s) 104 in the autonomous vehicle 700).

The sensor data pre-processor may use the sensor data 102 representative of one or more images (or other data representations) and load the sensor data into memory in the form of a multi-dimensional array/matrix (alternatively referred to as tensor, or more specifically an input tensor, in some examples). The array size may be computed and/or represented as  $W \times H \times C$ , where  $W$  stands for the image width in pixels,  $H$  stands for the height in pixels and  $C$  stands for the number of color channels. Without loss of generality, other types and orderings of input image components are also possible. Additionally, the batch size  $B$  may be used as a dimension (e.g., an additional fourth dimension) when batching is used. Batching may be used for training and/or for inference. Thus, the input tensor may represent an array of dimension  $W \times H \times C \times B$ . Any ordering of the dimensions may be possible, which may depend on the particular hardware and software used to implement the sensor data pre-processor. This ordering may be chosen to maximize training and/or inference performance of the machine learning model(s) 104.

A pre-processing image pipeline may be employed by the sensor data pre-processor to process a raw image(s) acquired by a sensor(s) and included in the sensor data 102 to produce the pre-processed sensor data which may represent an input image(s) to the input layer(s) (e.g., feature extractor layer(s) 126 of Fig. 1B) of the machine learning

model(s) 104. An example of a suitable pre-processing image pipeline may use a raw RCCB Bayer (e.g., 1-channel) type of image from the sensor and convert that image to an RCB (e.g., 3-channel) planar image stored in Fixed Precision (e.g., 16-bit-per-channel) format. The pre-processing image pipeline may include decompanding, noise reduction, demosaicing, white  
5 balancing, histogram computing, and/or adaptive global tone mapping (e.g., in that order, or in an alternative order).

Where noise reduction is employed by the sensor data pre-processor, it may include bilateral denoising in the Bayer domain. Where demosaicing is employed by the sensor data pre-processor, it may include bilinear interpolation. Where histogram computing  
10 is employed by the sensor data pre-processor, it may involve computing a histogram for the C channel, and may be merged with the decompanding or noise reduction in some examples. Where adaptive global tone mapping is employed by the sensor data pre-processor, it may include performing an adaptive gamma-log transform. This may include calculating a histogram, getting a mid-tone level, and/or estimating a maximum luminance with the mid-  
15 tone level.

The machine learning model(s) 104 may use as input one or more images (or other data representations) represented by the sensor data 102 to generate the boundary points 106 (e.g., representative of 2D pixel locations of boundary(ies) within the image) and/or the class labels 108 (e.g., boundary class labels corresponding to boundary classes associated  
20 with the boundary(ies)) as output. In a non-limiting example, the machine learning model(s) 104 may take as input an image(s) represented by the pre-processed sensor data and/or the sensor data 102, and may use the sensor data to regress on the boundary points 106 and to generate predictions for the class labels 108 that correspond to the boundary points 106.

Although examples are described herein with respect to using neural networks,  
25 and specifically convolutional neural networks, as the machine learning model(s) 104 (e.g., as described in more detail herein with respect to FIGs. 1B and 1C), this is not intended to be limiting. For example, and without limitation, the machine learning model(s) 104 described herein may include any type of machine learning model, such as a machine learning model(s) using linear regression, logistic regression, decision trees, support vector machines (SVM),  
30 Naïve Bayes, k-nearest neighbor (Knn), K means clustering, random forest, dimensionality reduction algorithms, gradient boosting algorithms, neural networks (e.g., auto-encoders, convolutional, recurrent, perceptrons, Long/Short Term Memory (LSTM), Hopfield,

Boltzmann, deep belief, deconvolutional, generative adversarial, liquid state machine, etc.), and/or other types of machine learning models.

The boundary points 106 may represent points (or pixels) of an image represented by the sensor data that correspond to one or more boundaries dividing drivable  
5 free-space from non-drivable space in a real-world environment of the vehicle 700. For example, for each boundary there may be one boundary point 106 (or more, in some examples) for each column of the image (e.g., as a result of regression analysis by the machine learning model(s) 104), which may depend on the spatial width of the image (e.g., for a 1920 W x 1020 H image, there may be 1920 boundary points 106 for each boundary of  
10 the one or more boundaries). As a result, a boundary may be determined in the real-world environment across the entire lateral dimension of the field of view of the sensor(s) (e.g., the camera(s)) based on the boundary points 106.

As described in more detail herein with respect to FIG. 1D, in some non-limiting examples the boundary points 106 may be output by the machine learning model(s)  
15 104 as a one dimensional array (e.g., with one row, and a number of columns corresponding to the spatial width of the image). In such examples, each cell may include a value between 0 and 1 corresponding to a percentage of the spatial height of the image where the boundary point 106 (or pixel) for that column is located (e.g., the boundary points 106 output by the machine learning model 104 may be normalized to the height of the image). As such, if the  
20 image has a spatial height of 100, and the value in the first cell corresponding to the first column of the image is 0.5, the boundary point 106 for the first column may be determined to be at pixel 50. A one dimensional array may be the output of the machine learning model(s) 104 where the machine learning model(s) is used to regress on the boundary within the image.

25 In other examples, the boundary points 106 output by the machine learning model(s) 104 may be in a different format, such as a value(s) corresponding to the actual pixel within the column, a value(s) based on a distance from center of the image, and/or another format. In such examples, the output may not be a one dimensional array, but may be of any dimension (e.g., a number of columns corresponding to the spatial width of the image  
30 and a number of rows corresponding to the spatial width of the image, and/or other dimensions).

The class labels 108 may represent, for each of the boundary points 106, a boundary class or type that corresponds to the boundary point 106. For example, for a first

boundary point 106 a class label 108 may be for a pedestrian, for a second boundary point 106 a class label 108 may be for a vehicle, for a third boundary point a class label 108 may be for a barrier, and so on. There may be any number of class labels 108 that the machine learning model(s) 104 is trained to identify and label. In one non-limiting example, there  
5 may be four class labels 108 corresponding to four classes including vehicles, pedestrians, barriers or curbs, and a catchall or unidentified class (e.g., boundary types that don't fall into any of the other boundary classes or cannot be confidently placed into one of the other boundary classes). In other examples, there may be less than four class labels 108 or more than four class labels 108, and/or the class labels 108 may be more or less granular. For  
10 example, the class labels 108 may correspond to less granular classes such as static or dynamic, or to more granular classes such as sedan, sport utility vehicle (SUV), truck, bus, adult, child, curb, divider, guardrail, etc.

As described in more detail herein with respect to FIG. 1D, in some non-limiting examples the class labels 108 may be output (e.g., computed) by the machine  
15 learning model(s) 104 as a two-dimensional array (e.g., with a number of rows corresponding to the number of classes the machine learning model(s) 104 is trained to identify and label, and a number of columns corresponding to the spatial width of the image). In such examples, each of the cells in each column of the array may include a value between 0 and 1 (e.g., with a sum of all values for each column equaling 1) corresponding to a  
20 probability or confidence of the class label 108 that corresponds to the boundary point 106 from the same column.

The boundary points 106 and/or the class labels 108 may undergo one or more post-processing operations using one or more post-processing techniques. For example, the boundary points 106 may undergo spatial smoothing 110. The spatial smoothing 110 may be  
25 executed with a Gaussian filter, as described herein, and may be applied to the boundary points 106 predicted (e.g., as the output) by the machine learning model(s) 104. The goal of the spatial smoothing 110 may be to eliminate abrupt changes in the boundary between adjacent sets of the boundary points 106. The spatial smoothing 110 may be executed on the boundary points 106 from a single frame of the sensor data in view of others of the boundary  
30 points 106 from the same single frame. For example, one or more of the boundary points 106 corresponding to the boundary from a single frame may be filtered or smoothed (e.g., using a Gaussian filter) by using information (e.g., values) corresponding to one or more adjacent boundary points 106. As such, when a first boundary point 106 is significantly different from

an adjacent second boundary point 106, the spatial smoothing 110 may be used to create a less harsh or drastic transition between the first and second boundary points 106.

In some non-limiting examples, the spatial smoothing 110 may be implemented using equation (1), below:

$$boundary_{spatial}[col_{curr.}] = \sum_{idx=0}^{length(gauss\_filter)} boundary_{curr.}[col_{curr.} - radius_{gauss\_filter} \div idx] * gauss\_filter[idx] \quad (1)$$

where *gauss\_filter* is a Gaussian filter array, *col\_curr.* is a current column, *radius\_gauss\_filter* is the radius of the Gaussian filter, *boundary\_spatial* is the boundary array after spatial smoothing, and *idx* is an index (e.g., that goes from 0 to a length of the Gaussian filter array in equation (1)).

As another example of post-processing, the boundary points 106 and/or the class labels 108 may undergo temporal smoothing 112. The goal of the temporal smoothing 112 may be to eliminate abrupt changes in the boundary points 106 and/or the class labels 108 between a current frame and one or more prior frames of the sensor data. As such, the temporal smoothing 112 may look at boundary points 106 and/or class labels 108 of a current frame in view of boundary points 106 and/or class labels 108 of one or more prior frames to smooth the values of the boundary points 106 and/or the class labels 108 so that differences between successive frames may not be as drastic. In some examples, the prior frames may each have been weighted or smoothed based on its prior frame(s), and as a result, each successive frame may be a weighted representation of a number of prior frames (e.g., all prior frames). As such, when a first boundary point 106 and/or a first class label 108 is significantly different from a prior second boundary point 106 and/or second class label 108, the temporal smoothing 112 may be used to create a less harsh or drastic transition between the first and second boundary points 106.

In some non-limiting examples, the temporal smoothing 112 for the boundary points 106 may be implemented using equation (2), below:

$$boundary_{smooth} = \alpha * boundary_{prev.} \div (1 - \alpha) * boundary_{curr.} \quad (2)$$

where  $\alpha$  is a real number between 0 and 1, *boundary\_prev.* denotes a boundary array (e.g., an array of boundary points 106) in a previous frame, *boundary\_curr.* denotes a boundary array in a current frame, *boundary\_smooth* denotes a boundary array after the temporal smoothing 112.

In some examples,  $\alpha$  may be set to a default value. For example, the default value may be 0.5 in one non-limiting example. However, the value for  $\alpha$  may be different depending on the embodiment. Where  $\alpha$  is 0.5, for example, the weighting between a boundary point 106 and an adjacent boundary point(s) 106 may be 50% for the boundary point 106 and 50% for the adjacent boundary point(s) 106. In examples where the boundary point 106 is to be weighted more heavily than the adjacent boundary point(s) 106, the value of  $\alpha$  may be lower, such as 0.3, 0.4, etc., and where the boundary points is to be weighted less heavily than the adjacent boundary point(s) 106, the value of  $\alpha$  may be higher, such as 0.6., 0.7, etc.

In some non-limiting examples, the temporal smoothing 112 for the class labels 108 may be implemented using equation (3), below:

$$label\_prob_{smooth} = \alpha * label\_prob_{prev.} + (1 - \alpha) * label\_prob_{curr.}$$

(3)

where  $\alpha$  is a real number between 0 and 1,  $label\_prob_{prev.}$  denotes a class label probability array (e.g., an array of predictions for class labels 108) in a previous frame,  $label\_prob_{curr.}$  denotes a class label probability array in a current frame,  $label\_prob_{smooth}$  denotes a class label probability array after the temporal smoothing 112.

In some examples,  $\alpha$  may be set to a default value. For example, the default value may be 0.5 in one non-limiting example. However, the value for  $\alpha$  may be different depending on the embodiment. Where  $\alpha$  is 0.5, for example, the weighting between a prior frame(s) and a current frame may be 50% for a current frame and 50% for prior frame(s). In examples where a current frame is weighted more heavily than a prior frame(s), the value of  $\alpha$  may be lower, such as 0.3, 0.4, etc., and where the current frame is to be weighted less heavily than the prior frame(s), the value of  $\alpha$  may be higher, such as 0.6., 0.7, etc.

The spatial smoothing 110 and/or the temporal smoothing 112 may individually be performed, may both be performed, and may be performed in any order. For example, only the spatial smoothing 110 or only the temporal smoothing 112 may be performed. As another example, both the spatial smoothing 110 and the temporal smoothing 112 may be performed. In such examples, the spatial smoothing 110 may be performed prior to the temporal smoothing 112, the temporal smoothing 112 may be performed prior to the spatial smoothing 110, and/or both may be performed simultaneously or in parallel. In some examples, the spatial smoothing 110 may be performed on the output after the temporal

smoothing 112, or the temporal smoothing 112 may be performed on the output after the spatial smoothing 110.

The boundary points 106 computed by the machine learning model(s) 104 – after post-processing in some examples – may be converted from 2D point or pixel locations of the sensor data (e.g., of an image) to 3D or 2D real-world coordinates. In some examples, location determiner 114 may determine a real-world, 3D location of the boundary based on the 2D point or pixel coordinates. Using intrinsic camera parameters (e.g., focal length,  $f$ , optical center ( $u_o, v_o$ ), pixel aspect ratio,  $\alpha$ , skew,  $s$ , etc.), extrinsic camera parameters (e.g., 3D rotation,  $R$ , translation,  $t$ , etc.), and/or a height of the camera with respect to a ground plane, a 3D distance from the boundary (e.g., the boundary delineated by the boundary points 106) to the camera center may be computed. The real-world coordinate system and the camera coordinate system may be assumed to be the same coordinate system. As a result, a projective matrix,  $P$ , may be expressed according to equation (4), below:

$$P = [K \ 0]_{3 \times 4} \quad (4)$$

where  $K$  is a 3 x 3 dimensional intrinsic camera matrix.

In order to perform 3D to 2D projection, for any point,  $[X, Y, Z]$ , on the road or driving surface, a corresponding pixel,  $[u, v]$ , may be determined in an image that satisfies equation (5), below:

$$\lambda \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = [K]_{3 \times 3} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \quad (5)$$

where  $\lambda$  is a scaling constant in homogeneous coordinates and  $K$  is a 3 x 3 dimensional intrinsic camera matrix.

A pixel,  $p$ , in the image may be mapped to an infinite number of 3D points along the line which connects the camera center and the pixel,  $p$ . As such, in order to determine the corresponding 3D position of the pixel,  $p$ , one or more constraints may be imposed and/or stereo vision may be leveraged, depending on the embodiment. In examples where a constraint is imposed, the constraint may be that the 3D point always lies on the ground plane of the road or driving surface. The ground plane, in some examples, may satisfy the relationship illustrated in equation (6), below:

$$n^T x + c = 0 \quad (6)$$

where  $x = [X \ Y \ Z]^T$ ,  $c$  is a constant related to a height of the camera center from the ground plane (e.g., if the camera center is 1.65 meters from the ground plane,  $c = 1.65$ ),  $T$  is time, and, when assuming that the camera is orthogonal to the ground plane, in the camera coordinate system  $n = [0 \ -1 \ 0]^T$ .

- 5                    Since the position of the pixel,  $p$ , on the boundary may be known, the 3D location corresponding to the pixel,  $p$ , may be derived from a 3D to 2D projection equation (7), as a function of  $\lambda$  and  $K$ :

$$x = [XY Z]^T = \lambda K^{-1} p \quad (7)$$

- 10                    where  $p = [u \ v \ 1]^T$ . Using the solution of the 3D point in the equation for the ground plane result in the following relationship, represented by equation (8), below:

$$\lambda n^T K^{-1} p + c = 0 \quad (8)$$

Thus,

$$\lambda = \frac{-c}{n^T K^{-1} p} \quad (9)$$

Combining equations (8) and (9), results in equation (10), below:

15                    
$$x = \frac{-c}{n^T K^{-1} p} (K^{-1} p) \quad (10)$$

The distance from the boundary to the vehicle 700 across the ground plane may be determined by subtracting the projection of the camera center on the ground plane from  $x$ , such as represented in equation (11), below:

$$d = \|x - C''\|_2 \quad (11)$$

- 20                    where  $C''$  is the projection of the camera center on the ground plane equaling  $h = [0 \ 1.65 \ 0]^T$ , assuming the height of the camera from the ground plane is 1.65 meters. In addition, a distance,  $d_f$ , may be determined from the boundary to the front of the vehicle 700 by considering the distance from the camera center to the front of the ground plane (e.g., along the z-axis), or to the front of the vehicle. For example, assuming that the camera center
- 25                    is 2.85 meters away from the front of the ground plane, the distance,  $d_f$ , may be determined using equation (12), below:

$$d_f = d - 2.85 \quad (12)$$

- Although this example of 3D to 2D projection for determining distance from the vehicle 700 to the boundary is described herein, this is not intended to be limiting. For
- 30                    example, any method of determining real-world coordinates from image coordinates may be used to determine the distance from the vehicle 700 to the boundary, and thus the drivable

free-space for the vehicle 700 to move around within. As an example, 3D object coordinates may be determined, a modeling transformation may be used to determine 3D real-world coordinates corresponding to the 3D object coordinates, a viewing transformation may be used to determine 3D camera coordinates based on the 3D real-world coordinates, a  
5 projection transformation may be used to determine 2D screen coordinates based on the 3D camera coordinates, and/or a window-to-viewport transformation may be used to determine 2D image coordinates from the 2D screen coordinates. In general, real-world coordinates may be converted to camera coordinates, camera coordinates may be converted to film coordinates, and film coordinates may be converted to pixel coordinates. Once the pixel  
10 coordinates are known, and the boundary is determined (e.g., the boundary points 106 are determined), the location of the boundary in the real-world coordinates may be determined using the known mapping from the boundary points 106 to the real-world coordinates.

In other examples, stereo vision may be leveraged to determine a correlation between the 2D and 3D coordinates, such as where two or more cameras are used. In such  
15 examples, stereo vision techniques may be used to determine distances to the boundary using information from two or more images generated from two or more cameras (or other sensors, such as LIDAR sensors).

The boundary points 106 (e.g., the real-world locations corresponding to the boundary points 106, after post-processing) and the class labels 108 (e.g., after post-  
20 processing) may be used by one or more layers of an autonomous driving software stack 146 (alternatively referred to herein as “drive stack 146”). The drive stack 146 may include a sensor manager (not shown), perception component(s) (e.g., corresponding to a perception layer of the drive stack 146), a world model manager 116, planning component(s) 118 (e.g., corresponding to a planning layer of the drive stack 146), control component(s) 120 (e.g.,  
25 corresponding to a control layer of the drive stack 146), obstacle avoidance component(s) 122 (e.g., corresponding to an obstacle or collision avoidance layer of the drive stack 146), actuation component(s) 124 (e.g., corresponding to an actuation layer of the drive stack 146), and/or other components corresponding to additional and/or alternative layers of the drive stack 146. The process 100 may, in some examples, be executed by the perception  
30 component(s), which may feed up the layers of the drive stack 146 to the world model manager, as described in more detail herein.

The sensor manager may manage and/or abstract sensor data 102 from the sensors of the vehicle 700. For example, and with reference to FIG. 7C, the sensor data 102

may be generated (e.g., perpetually, at intervals, based on certain conditions) by global navigation satellite system (GNSS) sensor(s) 758, RADAR sensor(s) 760, ultrasonic sensor(s) 762, LIDAR sensor(s) 764, inertial measurement unit (IMU) sensor(s) 766, microphone(s) 796, stereo camera(s) 768, wide-view camera(s) 770, infrared camera(s) 772, surround camera(s) 774, long range and/or mid-range camera(s) 798, and/or other sensor types.

The sensor manager may receive the sensor data from the sensors in different formats (e.g., sensors of the same type, such as LIDAR sensors, may output sensor data in different formats), and may be configured to convert the different formats to a uniform format (e.g., for each sensor of the same type). As a result, other components, features, and/or functionality of the autonomous vehicle 700 may use the uniform format, thereby simplifying processing of the sensor data. In some examples, the sensor manager may use a uniform format to apply control back to the sensors of the vehicle 700, such as to set frame rates or to perform gain control. The sensor manager may also update sensor packets or communications corresponding to the sensor data with timestamps to help inform processing of the sensor data by various components, features, and functionality of an autonomous vehicle control system.

A world model manager 116 may be used to generate, update, and/or define a world model. The world model manager 116 may use information generated by and received from the perception component(s) of the drive stack 146 (e.g., the locations of the boundary dividing drivable free-space from non-drivable space and the class labels 108). The perception component(s) may include an obstacle perceiver, a path perceiver, a wait perceiver, a map perceiver, and/or other perception component(s). For example, the world model may be defined, at least in part, based on affordances for obstacles, paths, and wait conditions that can be perceived in real-time or near real-time by the obstacle perceiver, the path perceiver, the wait perceiver, and/or the map perceiver. The world model manager 116 may continually update the world model based on newly generated and/or received inputs (e.g., data) from the obstacle perceiver, the path perceiver, the wait perceiver, the map perceiver, and/or other components of the autonomous vehicle control system.

The world model may be used to help inform planning component(s) 118, control component(s) 120, obstacle avoidance component(s) 122, and/or actuation component(s) 124 of the drive stack 146. The obstacle perceiver may perform obstacle perception that may be based on where the vehicle 700 is allowed to drive or is capable of driving (e.g., based on the location of the boundary determined from the boundary points 106

- 20 -

and/or based on the class labels 108), and how fast the vehicle 700 can drive without colliding with an obstacle (e.g., an object, such as a structure, entity, vehicle, etc.) that is sensed by the sensors of the vehicle 700.

The path perceiver may perform path perception, such as by perceiving  
5 nominal paths that are available in a particular situation. In some examples, the path perceiver may further take into account lane changes for path perception. A lane graph may represent the path or paths available to the vehicle 700, and may be as simple as a single path on a highway on-ramp. In some examples, the lane graph may include paths to a desired lane and/or may indicate available changes down the highway (or other road type), or may include  
10 nearby lanes, lane changes, forks, turns, cloverleaf interchanges, merges, and/or other information.

The wait perceiver may be responsible to determining constraints on the vehicle 700 as a result of rules, conventions, and/or practical considerations. For example, the rules, conventions, and/or practical considerations may be in relation to traffic lights,  
15 multi-way stops, yields, merges, toll booths, gates, police or other emergency personnel, road workers, stopped busses or other vehicles, one-way bridge arbitrations, ferry entrances, etc. In some examples, the wait perceiver may be responsible for determining longitudinal constraints on the vehicle 700 that require the vehicle to wait or slow down until some condition is true. In some examples, wait conditions arise from potential obstacles, such as  
20 crossing traffic in an intersection, that may not be perceivable by direct sensing by the obstacle perceiver, for example (e.g., by using sensor data 102 from the sensors, because the obstacles may be occluded from field of views of the sensors). As a result, the wait perceiver may provide situational awareness by resolving the danger of obstacles that are not always immediately perceivable through rules and conventions that can be perceived and/or learned.  
25 Thus, the wait perceiver may be leveraged to identify potential obstacles and implement one or more controls (e.g., slowing down, coming to a stop, etc.) that may not have been possible relying solely on the obstacle perceiver.

The map perceiver may include a mechanism by which behaviors are discerned, and in some examples, to determine specific examples of what conventions are  
30 applied at a particular locale. For example, the map perceiver may determine, from data representing prior drives or trips, that at a certain intersection there are no U-turns between certain hours, that an electronic sign showing directionality of lanes changes depending on the time of day, that two traffic lights in close proximity (e.g., barely offset from one another)

are associated with different roads, that in Rhode Island, the first car waiting to make a left turn at traffic light breaks the law by turning before oncoming traffic when the light turns green, and/or other information. The map perceiver may inform the vehicle 700 of static or stationary infrastructure objects and obstacles. The map perceiver may also generate  
5 information for the wait perceiver and/or the path perceiver, for example, such as to determine which light at an intersection has to be green for the vehicle 700 to take a particular path.

In some examples, information from the map perceiver may be sent, transmitted, and/or provided to server(s) (e.g., to a map manager of server(s) 778 of FIG.  
10 7D), and information from the server(s) may be sent, transmitted, and/or provided to the map perceiver and/or a localization manager of the vehicle 700. The map manager may include a cloud mapping application that is remotely located from the vehicle 700 and accessible by the vehicle 700 over one or more network(s). For example, the map perceiver and/or the localization manager of the vehicle 700 may communicate with the map manager and/or one  
15 or more other components or features of the server(s) to inform the map perceiver and/or the localization manager of past and present drives or trips of the vehicle 700, as well as past and present drives or trips of other vehicles. The map manager may provide mapping outputs (e.g., map data) that may be localized by the localization manager based on a particular location of the vehicle 700, and the localized mapping outputs may be used by the world  
20 model manager 116 to generate and/or update the world model.

The planning component(s) 118 may include a route planner, a lane planner, a behavior planner, and a behavior selector, among other components, features, and/or functionality. The route planner may use the information from the map perceiver, the map manager, and/or the localization manager, among other information, to generate a planned  
25 path that may consist of GNSS waypoints (e.g., GPS waypoints). The waypoints may be representative of a specific distance into the future for the vehicle 700, such as a number of city blocks, a number of kilometers, a number of feet, a number of miles, etc., that may be used as a target for the lane planner.

The lane planner may use the lane graph (e.g., the lane graph from the path  
30 perceiver), object poses within the lane graph (e.g., according to the localization manager), and/or a target point and direction at the distance into the future from the route planner as inputs. The target point and direction may be mapped to the best matching drivable point and direction in the lane graph (e.g., based on GNSS and/or compass direction). A graph search

algorithm may then be executed on the lane graph from a current edge in the lane graph to find the shortest path to the target point.

The behavior planner may determine the feasibility of basic behaviors of the vehicle 700, such as staying in the lane or changing lanes left or right, so that the feasible behaviors may be matched up with the most desired behaviors output from the lane planner. For example, if the desired behavior is determined to not be safe and/or available, a default behavior may be selected instead (e.g., default behavior may be to stay in lane when desired behavior or changing lanes is not safe).

The control component(s) 120 may follow a trajectory or path (lateral and longitudinal) that has been received from the behavior selector of the planning component(s) 118 as closely as possible and within the capabilities of the vehicle 700. In some examples, the remote operator may determine the trajectory or path, and may thus take the place of or augment the behavior selector. In such examples, the remote operator may provide controls that may be received by the control component(s) 120, and the control component(s) may follow the controls directly, may follow the controls as closely as possible within the capabilities of the vehicle, or may take the controls as a suggestion and determine, using one or more layers of the drive stack 146, whether the controls should be executed or whether other controls should be executed.

The control component(s) 120 may use tight feedback to handle unplanned events or behaviors that are not modeled and/or anything that causes discrepancies from the ideal (e.g., unexpected delay). In some examples, the control component(s) 120 may use a forward prediction model that takes control as an input variable, and produces predictions that may be compared with the desired state (e.g., compared with the desired lateral and longitudinal path requested by the planning component(s) 118). The control(s) that minimize discrepancy may be determined.

Although the planning component(s) 126 and the control component(s) 120 are illustrated separately, this is not intended to be limiting. For example, in some embodiments, the delineation between the planning component(s) 118 and the control component(s) 120 may not be precisely defined. As such, at least some of the components, features, and/or functionality attributed to the planning component(s) 118 may be associated with the control component(s) 120, and vice versa.

The obstacle avoidance component(s) 122 may aid the autonomous vehicle 700 in avoiding collisions with objects (e.g., moving and stationary objects). The obstacle

avoidance component(s) 122 may include a computational mechanism at a “primal level” of obstacle avoidance, and may act as a “survival brain” or “reptile brain” for the vehicle 700. In some examples, the obstacle avoidance component(s) 122 may be used independently of components, features, and/or functionality of the vehicle 700 that is required to obey traffic rules and drive courteously. In such examples, the obstacle avoidance component(s) may ignore traffic laws, rules of the road, and courteous driving norms in order to ensure that collisions do not occur between the vehicle 700 and any objects. As such, the obstacle avoidance layer may be a separate layer from the rules of the road layer, and the obstacle avoidance layer may ensure that the vehicle 700 is only performing safe actions from an obstacle avoidance standpoint. The rules of the road layer, on the other hand, may ensure that vehicle obeys traffic laws and conventions, and observes lawful and conventional right of way (as described herein).

In some examples, the locations of the boundary(ies) and the class labels 108 may be used by the obstacle avoidance component(s) 122 in determining controls or actions to take. For example, the drivable free-space may provide an indication to the obstacle avoidance component(s) 122 of where the vehicle 700 may maneuver without striking any objects, structures, and/or the like. As another example, such as where the class labels 108 include traversable and non-traversable labels, or where a first boundary is traversable and another boundary is not traversable, the obstacle avoidance component(s) 122 may – such as in an emergency situation – traverse the traversable boundary to enter potentially non-drivable space (e.g., a sidewalk, a grass area, etc.) that has been determined to be safer than continuing within the drivable free-space (e.g., as a result of a suspected collision).

In some examples, as described herein, the obstacle avoidance component(s) 122 may be implemented as a separate, discrete feature of the vehicle 700. For example, the obstacle avoidance component(s) 122 may operate separately (e.g., in parallel with, prior to, and/or after) the planning layer, the control layer, the actuation layer, and/or other layers of the drive stack 146.

Now referring to FIG. 1B, FIG. 1B is an illustration of an example machine learning model 104A for boundary identification, in accordance with some embodiments of the present disclosure. The machine learning model 104A may be one example of a machine learning model 104 that may be used in the process 100 of FIG. 1A. The machine learning model 104A may include or be referred to as a convolutional neural network and thus may

alternatively be referred to herein as convolutional neural network 104A or convolutional network 104A.

As described herein, the machine learning model 104A may use sensor data 102 (and/or pre-processed sensor data) (illustrated as an image in FIG. 1B) as an input. The  
5 sensor data 102 may include images representing image data generated by one or more cameras (e.g., one or more of the cameras described herein with respect to FIGs. 7A-7C). For example, the sensor data 102 may include image data representative of a field of view of the camera(s). More specifically, the sensor data 102 may include individual images generated by the camera(s), where image data representative of one or more of the individual  
10 images may be input into the convolutional network 104A at each iteration of the convolutional network 104A.

The sensor data 102 may be input as a single image, or may be input using batching, such as mini-batching. For example, two or more images may be used as inputs together (e.g., at the same time). The two or more images may be from two or more sensors  
15 (e.g., two or more cameras) that captured the images at the same time.

The sensor data 102 and/or pre-processed sensor data may be input into a feature extractor layer(s) 126 of the convolutional network 104 (e.g., feature extractor layer 126A). The feature extractor layer(s) 126 may include any number of layers 126, such as the layers 126A-126C. One or more of the layers 126 may include an input layer. The input  
20 layer may hold values associated with the sensor data 102 and/or pre-processed sensor data. For example, when the sensor data 102 is an image(s), the input layer may hold values representative of the raw pixel values of the image(s) as a volume (e.g., a width, W, a height, H, and color channels, C (e.g., RGB), such as 32 x 32 x 3), and/or a batch size, B (e.g., where batching is used)

25 One or more layers 126 may include convolutional layers. The convolutional layers may compute the output of neurons that are connected to local regions in an input layer (e.g., the input layer), each neuron computing a dot product between their weights and a small region they are connected to in the input volume. A result of a convolutional layer may be another volume, with one of the dimensions based on the number of filters applied (e.g.,  
30 the width, the height, and the number of filters, such as 32 x 32 x 12, if 12 were the number of filters).

One or more of the layers 126 may include a rectified linear unit (ReLU) layer. The ReLU layer(s) may apply an elementwise activation function, such as the max (0,

x), thresholding at zero, for example. The resulting volume of a ReLU layer may be the same as the volume of the input of the ReLU layer.

One or more of the layers 126 may include a pooling layer. The pooling layer may perform a down-sampling operation along the spatial dimensions (e.g., the height and the width), which may result in a smaller volume than the input of the pooling layer (e.g., 16 x 16 x 12 from the 32 x 32 x 12 input volume). In some examples, the convolutional network 104A may not include any pooling layers. In such examples, strided convolution layers may be used in place of pooling layers. In some examples, the feature extractor layer(s) 126 may include alternating convolutional layers and pooling layers.

One or more of the layers 126 may include a fully connected layer. Each neuron in the fully connected layer(s) may be connected to each of the neurons in the previous volume. The fully connected layer may compute class scores, and the resulting volume may be 1 x 1 x number of classes. In some examples, the feature extractor layer(s) 126 may include a fully connected layer, while in other examples, the fully connected layer of the convolutional network 104A may be the fully connected layer separate from the feature extractor layer(s) 126. In some example, no fully connected layers may be used by the feature extractor 126 and/or the machine learning model 104A as a whole, in an effort to increase processing times and reduce computing resource requirements. In such examples, where no fully connected layers are used, the machine learning model 104A may be referred to as a fully convolutional network.

One or more of the layers 126 may, in some examples, include deconvolutional layer(s). However, the use of the term deconvolutional may be misleading and is not intended to be limiting. For example, the deconvolutional layer(s) may alternatively be referred to as transposed convolutional layers or fractionally strided convolutional layers. The deconvolutional layer(s) may be used to perform up-sampling on the output of a prior layer. For example, the deconvolutional layer(s) may be used to up-sample to a spatial resolution that is equal to the spatial resolution of the input images (e.g., the sensor data 102) to the convolutional network 104B, or used to up-sample to the input spatial resolution of a next layer.

Although input layers, convolutional layers, pooling layers, ReLU layers, deconvolutional layers, and fully connected layers are discussed herein with respect to the feature extractor layer(s) 126, this is not intended to be limiting. For example, additional or

alternative layers 126 may be used in the feature extractor layer(s) 126, such as normalization layers, SoftMax layers, and/or other layer types.

The output of the feature extractor layer(s) 126 may be an input to boundary point layer(s) 128 and/or class label layer(s) 130. The boundary point layer(s) 128 and/or the class label layer(s) 130 may use one or more of the layer types described herein with respect to the feature extractor layer(s) 126. As described herein, the boundary point layer(s) 128 and/or the class label layer(s) 130 may not include any fully connected layers, in some examples, to reduce processing speeds and decrease computing resource requirements. In such examples, the boundary point layer(s) 128 and/or the class label layer(s) 130 may be referred to as fully convolutional layers.

Different orders and numbers of the layers 126, 128, and 130 of the convolutional network 104A may be used depending on the embodiment. For example, for a first vehicle, there may be a first order and number of layers 126, 128, and/or 130, whereas there may be a different order and number of layers 126, 128, and/or 130 for a second vehicle; for a first camera, there may be a different order and number of layers 126, 128, and/or 130 than the order and number of layers for a second camera. In other words, the order and number of layers 126, 128, and/or 130 of the convolutional network 104A is not limited to any one architecture.

In addition, some of the layers 126, 128, and/or 130 may include parameters (e.g., weights and/or biases) – such as the feature extractor layer(s) 126, the boundary point layer(s) 128, and/or the class label layer(s) 130 – while others may not, such as the ReLU layers and pooling layers, for example. In some examples, the parameters may be learned by the machine learning model(s) 104A during training. Further, some of the layers 126, 128, and/or 130 may include additional hyper-parameters (e.g., learning rate, stride, epochs, kernel size, number of filters, type of pooling for pooling layers, etc.) – such as the convolutional layer(s), the deconvolutional layer(s), and the pooling layer(s) – while other layers may not, such as the ReLU layer(s). Various activation functions may be used, including but not limited to, ReLU, leaky ReLU, sigmoid, hyperbolic tangent (tanh), exponential linear unit (ELU), etc. The parameters, hyper-parameters, and/or activation functions are not to be limited and may differ depending on the embodiment.

In any example, the output of the machine learning model 104A may be the boundary points 106 and/or the class labels 108. In some examples, the boundary point layer(s) 128 may output the boundary points 106 and the class label layer(s) 130 may output

the class labels 108. As such, the feature extractor layer(s) 126 may be referred to as a first convolutional stream, the boundary point layer(s) 128 may be referred to as a second convolutional stream, and/or the class label layer(s) 130 may be referred to as a third convolutional stream.

5                In some examples, such as where there are two or more boundaries regressed on by the machine learning model(s) 104, for each boundary, there may be a separate stream of boundary point layer(s) and/or class label layer(s) (not shown). For example, for a second boundary in addition to the first boundary, a fourth convolutional stream of boundary point layer(s) may be used to determine the boundary points 106 for the second boundary, and a  
10                fifth convolutional stream of class label layer(s) 130 may be used to determine the associated class labels 108 for the second boundary. As such, for each boundary, there may be associated boundary point layer(s) and/or class label layer(s). In such examples, an additional class, or an additional parameter, may be learned by the machine learning model(s) 104 for traversable or non-traversable boundaries. For example, some boundaries separating  
15                drivable free-space from non-drivable space may be traversable, such as a curb, a grass or gravel area, etc., while other boundaries may be non-traversable, such as building, vehicle, pedestrians, etc. As such, the machine learning model(s) 104 may be trained to determine parameters associated with boundary classes or types, such as traversable and non-traversable boundary classes or types, in addition to, or alternatively from, determining the classes  
20                themselves (e.g., vehicles, pedestrians, curbs, barrier, etc.).

                 Now referring to FIG. 1C, FIG. 1C is another illustration of an example machine learning model 104B for boundary identification, in accordance with some embodiments of the present disclosure. The machine learning model 104B may be a non-limiting example of the machine learning model(s) 104 for use in the process 100 of FIG. 1A.  
25                The machine learning model 104B may be a convolutional neural network and thus may be referred to herein as convolutional neural network 104B or convolutional network 104B. In some examples, the convolutional network 104B may include any number and type of different layers, although some examples do not include any fully connected layers in order to increase processing speeds and reduce computing requirements to enable the process 100  
30                to run in real-time (e.g., at 30 fps or greater).

                 The convolutional network 104B may include feature extractor layer(s) 132, boundary point layer(s) 134, and/or class label layer(s) 136, which may correspond to the feature extractor layer(s) 126, the boundary point layer(s) 128, and/or the class label layer(s)

130 of FIG. 1B, respectively, in some examples. The feature extractor layer(s) 132 may include any number of layers, however, in some examples, the feature extractor layers 132 include eighteen or less layers or less layers in order to minimize data storage requirements and to increase processing speeds for the convolutional network 104B. In some examples, 5 the feature extractor layer(s) 132 includes convolutional layers that use 3 x 3 convolutions for each of its layers, with the exception of the first convolutional layer, in some examples, which may use a 7 x 7 convolutional kernel. In addition, in some example, the feature extractor layer(s) 132 may not include any skip-connections, which differs from conventional systems and may increase the processing times and accuracy of the system.

10 In some examples, the feature extractor layer(s) 132 may be similar to the structure illustrated in FIG. 8, and described in the accompanying text, of U.S. Provisional Patent Application Number 62/631,781, entitled "Method for Accurate Real-Time Object Detection and for Determining Confidence of Object Detection Suitable for Autonomous Vehicles", filed February 18, 2018 (hereinafter the '781 Application). However, in some 15 examples, the feature extractor layer(s) 132 of the present disclosure may include a network stride of 32, as compared to 16 in the structure of the '781 Application (e.g., the input may be down-sampled by 32 instead of 16). By using 16, rather than 32 as the stride, the convolutional network 104B may be computationally faster while not losing much, if any, accuracy. In addition, in some examples, the feature extractor layer(s) 132 may use an 20 applied pool size of 2 x 2, rather than 3 x 3 as disclosed in the '781 Application.

The feature extractor layer(s) 132 may continuously down sample the spatial resolution of the input image until the output layers are reached (e.g., down-sampling from a 960 x 544 x 3 input spatial resolution to the feature extractor layer(s) 132, to 480 x 272 x 64 at the output of the first feature extractor layer, to 240 x 136 x 64 at the output of the second 25 feature extractor layer, to 30 x 17 x 12 as output of the last of the feature extractor layer(s) 132). The feature extractor layer(s) 132 may be trained to generate a hierarchical representation of the input image(s) (or other sensor data representations) received from the sensor data 102 and/or pre-processed sensor data with each layer generating a higher-level extraction than its preceding layer. In other words, the input resolution across the feature 30 extractor layer(s) 132 (and/or any additional or alternative layers) may be decreased, allowing the convolutional network 104B to be capable of processing images faster than conventional systems.

The boundary point layer(s) 134 and/or the class label layer(s) 136 may take the output of the feature extractor layer(s) 132 as input. The boundary point layer(s) 134 may be used to regress on the boundary points 106 and the class label layer(s) 136 may be used to predict class labels 108 that may correspond to the boundary points 106 regressed on by the  
5 boundary point layer(s) 132. The boundary point layer(s) 134 and/or the class label layer(s) 136 may include any number or type of layers, but in some examples, the boundary point layers 134 and/or the class label layer(s) 136 may include three layers.

For example, for the boundary point layer(s) 134, a first layer, Conv. 128, may reduce the number of feature maps from the output of the feature extractor layer(s) 132. A  
10 second layer, Deconv. 1, may up-sample the input from the first layer along the width, in order to get one point along each column of the image (e.g., based on the spatial width, and thus the number of columns of pixels, of the input to the machine learning model 104B, such as the sensor data 102). The third or output layer, Conv. 1, may reduce the height of the output array to one.

15 For the class label layers 136, a first layer, Conv. 128, may reduce the number of feature maps from the output of the feature extractor layer(s) 132. A second layer, Deconv. 4 may up-sample the input from the first layer along the width, in order to get a number of class label predictions along each column of the image that corresponds to the number of classes the machine learning model 104B was trained to predict (e.g., in this  
20 example, the number of classes may be four). The third or output layer, Conv. 4, may reduce the height of the output array to the number of classes the machine learning model 104B is trained to predict (e.g., four) such that for each column, the output array includes a confidence or probability for each boundary class or type for the corresponding boundary point 106 from the corresponding column in the output of the boundary point layers 134.

25 In some examples, with reference to FIG. 1D, the output of the machine learning model(s) 104 may include the boundary points 106 and the class label(s) 108. For example, the output of the boundary points 106, as described herein, may include a one dimensional array having one row 138 and a number (1-N) of columns 140 corresponding to the spatial width of the input image to the machine learning model(s) 104. For example, if  
30 the input image has a spatial width dimension of 1080 pixels, then N may be 1080. As such, each cell of the array may include a value (from 0 to 1) that corresponds to a percentage of the spatial height of the input image where the boundary point 106 is located (e.g., the pixel that corresponds to the first boundary in the column, from bottom to top, that separates

drivable free-space from non-drivable space). As an example, column 140(1) may correspond to the first column of pixels in the input image, and may have a value of 0.4 (e.g., 40% of the height). As such, if the spatial height dimension of the image were 920 pixels, then the 368<sup>th</sup> pixel (e.g., 40% of 920) from the bottom of the column 140(1) may be identified as the boundary point 106 for the column 140(1) of the input image.

As another example, the output of the class labels 108, as described herein, may include a two dimensional array having a number (1-K) of rows 142 corresponding to a number of boundary classes or types the machine learning model(s) 104 is trained to predict. As such, each row may correspond to a different boundary class or type. The number (1-M) of columns 144 may correspond to the spatial width of the input image to the machine learning model(s) 104. For example, if the machine learning model(s) 104 is trained to predict five classes, then K may be 5, and if the input image has a spatial width dimension of 1080 pixels, then M may be 1080.

As such, each cell of the array may include a value (from 0 to 1) that corresponds to a confidence or probability of the boundary point 106 of the corresponding column 140 in the boundary point array belonging to the boundary class or type associated with the row 142. As an example, column 144(1) may correspond to column 140(1) of the boundary point array, and row 142(1) may correspond to a class label 108 for pedestrians. As such, the value, 0.7, of the cell at row 142(1) and column 144(1) may correspond to a confidence or percentage that the boundary class or type of the boundary point of row 138, column 140(1) of the boundary point array is a pedestrian boundary class or type. In this example, assuming that 0.7 is the highest value, the determination may be that the boundary point 106 at row 138, column 140(1) of the boundary point array corresponds to a pedestrian.

This process or correlating boundary point locations with class labels may be performed for each of the columns of the input image. As a result, and after post-processing in some examples, the drive stack 146 may determine one or more operations for navigating or controlling the vehicle 700 based on the pixel locations of the boundary points 106 in the input image and the class labels 108 determined to be associated therewith.

Now referring to FIG. 2, each block of method 200, and other methods described herein, comprises a computing process that may be performed using any combination of hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory. The methods may also be embodied as computer-usable instructions stored on computer storage media. The

methods may be provided by a standalone application, a service or hosted service (standalone or in combination with another hosted service), or a plug-in to another product, to name a few. In addition, method 200 is described, by way of example, with respect to the process 100 of FIG. 1A. However, the method 200 may additionally or alternatively be executed by  
5 any one system, or any combination of systems, including, but not limited to, those described herein.

FIG. 2 is an example flow diagram for a method 200 of determining a boundary, in accordance with some embodiments of the present disclosure. The method 200, at block  
10 B202, includes receiving sensor data. For example, the sensor data 102 and/or pre-processed sensor data may be received. The sensor data may be captured by one or more sensors of the vehicle 700 and may represent a field(s) of view of the sensor(s), or may otherwise represent one or more aspect of the physical environment (e.g., for a GPS sensor, the sensor data may represent coordinates or locations within the physical environment). For example, the sensor  
15 data 102 may include image data representative of image(s) of a field(s) of view of one or more cameras of the vehicle 700.

The method 200, at block B204, includes applying the sensor data to a neural network. For example, the sensor data and/or the pre-processed sensor data may be applied to the machine learning model(s) 104. The machine learning model(s) 104, as described  
20 herein, may be trained to compute the boundary points 106 and the class labels 108.

The method 200, at block B206, includes computing, by a first stream of the neural network, a set of points that correspond to a boundary driving drivable free-space from non-drivable space. For example, the machine learning model(s) 104 may compute, based on the sensor data, the boundary points 106 for one or more boundaries. The boundary points  
25 106 may be computed by a first stream (e.g., the boundary point layers 128) of the machine learning model(s) 104.

The method 200, at block B208, includes computing, by a second stream of the neural network, class labels for points of the set of points. For example, the machine learning model(s) 104 may compute, based on the sensor data, class labels 108 that  
30 correspond to the boundary points 106. The class labels 108 may be computed by a second stream (e.g., the class label layers 130) of the machine learning model(s) 104.

The method 200, at block B210, includes determining locations within a physical environment corresponding to the sets of points. For example, the point or pixel

locations determined from the boundary points 106 may be used to determine corresponding locations within the real-world environment for the boundary dividing drivable free-space from non-drivable space.

The method 200, at block B212, includes executing one or more operations for  
5 controlling the vehicle. For example, one or more operations may be executed for controlling the vehicle 700 within the real-world environment. In some examples, the class labels and real-world locations determined from the locations of the boundary points 106 may be used by the drive stack 146 to determine the one or more operations (e.g., world model updates, trajectories, controls, actuations, etc.).

10 The method 200 may repeat, as indicated by the dashed arrow, for each sensor data input.

### **Training the Machine Learning Model(s)**

In some examples, the machine learning model(s) 104 (e.g., a neural  
network(s)) may be trained with training images (e.g., training images 300A, 300B, and/or  
15 500) that include an image and one or more associated annotations, labels, or masks. For example, annotations may be associated with the image and the training images and may be used as ground truth data for training the machine learning model(s) 104. The annotations for the training images may be generated within a drawing program (e.g., an annotation program), a computer aided design (CAD) program, a labeling program, another type of  
20 program suitable for generating the annotations, and/or may be hand drawn, in some examples. In any example, the annotations may be synthetically produced (e.g., generated from computer models or renderings), real produced (e.g., designed and produced from real-world data), machine-automated (e.g., using feature analysis and learning to extract features from data and then generate labels), human annotated (e.g., labeler, or annotation expert,  
25 defines the location of the labels), and/or a combination thereof (e.g., human formulates one or more rules or labeling conventions, machine generates annotations).

The outputs of the machine learning model(s) 104 may be compared against the ground truth data (e.g., the training images) using one or more loss functions. For example, the machine learning model(s) 104 may be trained with the training images using  
30 multiple iterations until the value of a loss function(s) of the machine learning model(s) 104 is below a threshold loss value (e.g., acceptable loss).

For example, the machine learning model(s) 104 may perform forward pass computations on the representations (e.g., image(s)) of the sensor data 102 and/or pre-

processed sensor data to generate the boundary points 106 and the class labels 108. The loss function(s) may be used to measure error in the predictions of the machine learning model(s) 104 using ground truth data, as described in more detail herein with respect to at least FIGs. 3A-6. In some examples, a cross entropy loss function (e.g., binary cross entropy) and/or an L1 loss function may be used.

In some examples, different loss functions may be used for the boundary points 106 and the class labels 108. For example, the boundary points 106 may be associated with an L1 loss function, represented by the relationship of equation (13), below:

$$\text{L1 Loss: } L(y, x) = |y - f_w(x)| \quad (13)$$

where  $y$  is an array of ground truth boundary points based on the annotations,  $|$  is the absolute value function, and  $f_w$  is an array of predicted boundary points 106 output by the machine learning model(s) 104 (e.g., the boundary point array of FIG. 1D).

As another example, the class labels 108 may be associated with a cross entropy loss function, represented by the relationship of equation (14), below:

$$\text{Cross Entropy Loss: } L(y, x) = -\sum_{i=1}^n y_i \log(f_{i,w}(x)) \quad (14)$$

where  $y_i$  is a binary indicator (1 or 0) if  $i$  is the class label for the input,  $x$ ,  $f_{i,w}$  is the probability that the machine learning model(s) 104 predicts that  $i$  is the class label for the input,  $x$ , and  $\log$  denotes the logarithmic function.

In some examples, a weighted combination of both the boundary point loss function (e.g., L1 loss) and the class label loss function (e.g., cross entropy loss) may be used to compute the final loss. For example, a weight ratio of 1:10 (e.g., boundary point: class label), may be used. However, this is not intended to be limiting any other weighting ratios may be used without departing from the scope of the present disclosure. In some examples, an auto-weight scheme may be used that adjusts cost weights adaptively or automatically when given only the target weight ratios for the weighted losses.

For example, a boundary point cost and corresponding cost weight may be denoted by  $L_{\text{boundary}}$  and  $w_{\text{boundary}}$ , respectively, and the class label cost and corresponding cost weight may be denoted by  $L_{\text{label}}$  and  $w_{\text{label}}$ , respectively. In such an example, the weights,  $w$ , may be selected such that the relationship of equation (15) is satisfied:

$$\frac{L_{\text{boundary}} w_{\text{boundary}}}{L_{\text{label}} w_{\text{label}}} = K_{\text{boundary}} \quad (15)$$

- 34 -

where  $K_{boundary}$  is the target ratio for the weighted loss. As a result, since the weights also need to have a sum of one, the follow expressions of equations (16) and (17) may be derived:

$$W_{boundary} = \frac{K_{boundary} L_{label}}{K_{boundary} L_{label} + L_{boundary}} \quad (16)$$

$$W_{label} = \frac{L_{boundary}}{K_{boundary} L_{label} + W_{boundary}} \quad (17)$$

This rule may be used periodically, such as once per training epoch, to set the weights. As such, the costs in the rule may be calculated from the training data (e.g., the ground truth data) to prevent overfitting to the training data. The weights may thus be set adaptively in a way that, after convergence, the desired target ratio for the weighted losses may be achieved.

Backward pass computations may be performed to recursively compute gradients of the loss function with respect to training parameters. In some examples, weights and biases of the machine learning model(s) 104 may be used to compute these gradients. In some examples, an optimizer may be used to make adjustments to the training parameters (e.g., weights, biases, etc.). For example, one or more of an Adam optimizer, stochastic gradient descent, or stochastic gradient descent with a momentum term, may be used. The training process may be reiterated until the trained parameters converge to optimum, desired, and/or acceptable values.

Now referring to FIG. 3A, FIG. 3A is an example training image 300A for a machine learning model, in accordance with some embodiments of the present disclosure. The annotations associated with the image may include annotating, along each column of pixels of the image, from bottom to top of the image, one point corresponding to a first location of a boundary dividing drivable free-space from non-drivable space. As a result, once a point is identified and labeled in each column, a full boundary may be labeled. The annotations may include associated class labels such that each point along the boundary has an associated boundary class or type (e.g., pedestrian, vehicle, curb, barrier, catchall, etc.). In some examples, the annotations may include segmented or discontinuous boundaries (e.g., so long as one pixel for each column of pixels of the image includes an annotation), while in other examples, the annotations may include continuous boundary curves (e.g., the segments may be connected to form a continuous curve).

With respect to the training image 300A, and assuming the machine learning model(s) 104 is trained to predict at least three classes (e.g., pedestrians, vehicles, and/or curbs or barriers), the training image 300A may include associated annotations 312 for a curb or barrier 306 class label (e.g., as indicated by the curb or barrier 306A and 306B), annotations 310 for a vehicle 304 class label (e.g., as indicated by the vehicles 304A and 304B), and annotations 308 for a pedestrian 302 class label (e.g., as indicated by the pedestrians 302A and 302B). In some examples, the annotations may conform to the shape or contour of the boundary (e.g., the annotations 312 for the barriers 306 and the annotations 310 for the vehicle 304), while in other examples, such as where the boundary class or type is more difficult to delineate, the annotations may not conform to the shape or contour of the boundary (e.g., the annotations 308 for the pedestrians 302). Where the annotations do not conform to the shape or contour of the boundary, line segments (e.g., curved line segments, straight line segments, etc.), may be used. In the example of the pedestrians 302, the annotations 308 may include straight line segments that extend along the lateral dimension of the pedestrians from a bottom most portion of the pedestrian 302 (e.g., from a point closest to the bottom of the image and thus closest to the vehicle 700 in real-world coordinates). For example, even though an annotation following a gap between legs of the pedestrian 302A may most closely conform to the shape of the pedestrian 302A, the straight line annotation 308 may be used to ensure safety and accuracy. More specifically, the annotations 308, or other similar annotations for other boundary types (e.g., vehicles, in some examples), may be done in order to prevent any space occupied by any portion of the pedestrian 302A from being determined to be drivable free-space. Especially with respect to dynamic objects, undershooting the predicted drivable free-space may be preferable to overshooting.

In some examples, such as where a road or driving surface extends into a horizon, the annotations and/or the training image may be truncated or clamped. For example, a percentage of the top of the image and/or the associated annotations may be removed, and/or the machine learning model(s) 104 may be trained to ignore, or give less weight to, the upper clamped or truncated portions of the image. As a result, the drive stack 146 may not take into account, or may take into account to a lesser degree, the determined boundary points 106 and/or class labels 108 at further distances.

FIG. 3B is another example training image 300B for a machine learning model, in accordance with some embodiments of the present disclosure. With respect to the training image 300B, and assuming the machine learning model(s) 104 is trained to predict at

- 36 -

least three classes (e.g., vehicles, curbs or barriers, and a catchall class), the training image 300B may include associated annotations 314 for a curb or barrier 306 class label (e.g., as indicated by the curb or barrier 306C and 306D), annotations 316 for a vehicle 304 class label (e.g., as indicated by the vehicles 304C and 304D), and annotations 318 for a catchall class label. In some examples, such as in the training image 300B, an attempt may be made to train the machine learning model(s) 104 to avoid sharp or erratic changes between boundaries. As a result, the boundary points may be annotated such that the sharp transitions are avoided. With respect to the training image 300B, the boundary points between the vehicle 304D and the curb or barrier 306C, and the boundary points between the vehicle 304D and the vehicle 304C may have a smoother transition, as illustrated. In such examples, the portion of the boundary (e.g., corresponding to the annotations 318) that transitions between the vehicle 304D and the curb or barrier 306C and the portion of the boundary that transitions between the vehicle 304D and the vehicle 304C may be labeled as a catchall class (or may be labeled as a vehicle class to conform with the dynamic nature of the surrounding vehicles 304C and 304D). In any example, by training the machine learning model(s) 104 to determine less sharp or drastic changes between boundary types, and thus to determine a more smooth boundary output, the boundary dividing drivable free-space from non-drivable space may be more clearly understood by the vehicle 700 (e.g., by the drive stack 146).

Now referring to FIG. 4, each block of method 400, described herein, comprises a computing process that may be performed using any combination of hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory. The method may also be embodied as computer-usable instructions stored on computer storage media. The method may be provided by a standalone application, a service or hosted service (standalone or in combination with another hosted service), or a plug-in to another product, to name a few. In addition, method 400 is described, by way of example, with respect to the process 100 of FIG. 1A. However, this method may additionally or alternatively be executed by any one system, or any combination of systems, including, but not limited to, those described herein.

FIG. 4 is an example flow diagram for a method 400 of training a machine learning model, in accordance with some embodiments of the present disclosure. The method 400, at block B402, includes receiving a training image. For example, a training image may be received that may have been generated by a vehicle (e.g., the vehicle 700 or another vehicle) or that may have been generated from a virtual environment.

The method 400, at block B404, includes receiving annotation data representative of one or more annotations corresponding to the training data. For example, annotation data may be received that represents one or more boundary labels that may correspond to annotated boundary points. In addition, annotation data may be received that  
5 represents class labels. For example, before, during, or after the annotating of the boundary points, class labels may be associated with the annotations of the boundary labels. In some examples, a different tool within a labeling program may be used depending on the boundary class or type that the annotated boundary label corresponds to. In such examples, the annotated boundary label may represent or be a contextual or semantic representation of the  
10 boundary class or type. The result of the annotations may be visually represented similar to the training images 300A and 300B.

The method 400, at block B406, includes training a neural network using the annotations as ground truth data. For example, the training images including the associated annotations may be used to train the machine learning model(s) 104 (e.g., using the loss  
15 functions described herein).

Now referring to FIG. 5, FIG. 5 is another example training image 500 for a machine learning model, in accordance with some embodiments of the present disclosure. The training image 500 may include associated annotations for training the machine learning model(s) to predict two or more boundaries and associated class labels. For example, a first  
20 annotation 502 (indicated at least in part by segments 502A, 502B, and 502C) may correspond to a first boundary and a second annotation 504 (indicated at least in part by segments 504A, 504B, and 504C) may correspond to a second boundary may be generated for association with the image. In some examples, the machine learning model(s) 104 may be trained to generate additional boundaries, from top to bottom of the image, until a non-  
25 traversable boundary is determined, or until a threshold amount of the spatial width of the image includes non-traversable boundary segments (e.g., that may be a combination of segments from two or more different boundaries). In such examples, there may be additional annotations for each training image 500 corresponding to each of the boundaries. In other examples, the machine learning model(s) 104 may be trained to generate a set number of  
30 boundaries, such as two boundaries with respect to FIG. 5. In any example, the boundaries may be annotated similar to the training images 300A and 300B, described herein, and the training images 300A and 300B may be annotated similar to the training image 500.

The training image 500 may include the annotation 502 of the first boundary and the annotation 504 of the second boundary. Assuming that the machine learning model(s) 104 is trained to detect at least four classes and/or parameters (e.g., vehicles, curbs or barriers, traversable, non-traversable), the annotations 502 of the first boundary may include a curb or barrier 516 class label, a vehicle 506 class label, a class label or parameter for a non-traversable boundary type (e.g., the vehicle 506, along segment 502B), and/or a class label or parameter for a traversable boundary type (e.g., the curb or barrier 516, along segments 502A). Similarly, with respect to the second boundary, the annotations 504 may include a curb or barrier 514, 510, 508 class label, a vehicle 512 class label, a class label or parameter for a non-traversable boundary type (e.g., the vehicle 512, the curb or barrier 510 and 508, along segments 504A and 504C), and/or a class label or parameter for a traversable boundary type (e.g., the curb or barrier 514, along segments 504B).

In some examples, depending on a length of a segment of a boundary that is traversable (e.g., corresponding to the segment 504C of the annotation 504), the segment may be annotated as non-traversable. For example, even though the type of boundary (e.g., a curb) may be traversable, the surrounding information may be used to determine if the segment is actually traversable. In this example, the segment 504C may not be traversable, or may not safely be traversable, because a vehicle 512 and a barrier 510 may be immediately adjacent the segment 504C. As such, the segment 504C may be annotated as non-traversable.

In some examples, such as in the illustration of FIG. 5, there may be two or more drivable free-spaces. For example, there may be a first drivable free-space 520 and second drivable free-space 518. In such examples, the curb or barrier 516 may include a boundary separating the first drivable free-space 520 from the second drivable free-space 518. In some examples, there may be an additional boundary (not shown) identified on an opposite side of the curb or barrier 516 (e.g., opposite the boundary identified by the annotation 502). In such examples, the annotation 502 may correspond to a boundary dividing drivable free-space from non-drivable space, and the additional boundary may divide the non-drivable space from the second drivable free-space 518.

As another example, such as where a lane merge is captured in the sensor data, the area between the two-lanes prior to the merge may include non-drivable space. In such an example, the annotations may include annotating a first boundary dividing drivable free-space of a first lane from the non-drivable space of the area between the two lanes prior to merge, and a second boundary dividing the non-drivable space of the area from the drivable

free-space of a second lane (e.g., the lane merging into the first lane). As such, the machine learning model(s) 104 may be trained to annotate more than one drivable free-space and/or the non-drivable space therebetween.

Now referring to FIG. 6, each block of method 600, described herein,  
5 comprises a computing process that may be performed using any combination of hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory. The method may also be embodied as computer-usable instructions stored on computer storage media. The method may be provided by a standalone application, a service or hosted service (standalone or in combination with another  
10 hosted service), or a plug-in to another product, to name a few. In addition, method 600 is described, by way of example, with respect to the process 100 of FIG. 1A. However, this method may additionally or alternatively be executed by any one system, or any combination of systems, including, but not limited to, those described herein.

FIG. 6 is an example flow diagram for a method 600 for multiple boundary  
15 detection, in accordance with some embodiments of the present disclosure. The method 600, at block B602, includes receiving sensor data. For example, the sensor data 102 and/or pre-processed sensor data may be received. The sensor data may be captured by one or more sensors of the vehicle 700 and may represent a field(s) of view of the sensor(s). For example, the sensor data 102 may include image data representative of image(s) of a field(s) of view of  
20 one or more cameras of the vehicle 700.

The method 600, at block B604, includes applying the sensor data to a machine learning model(s). For example, the sensor data and/or the pre-processed sensor data may be applied to the machine learning model(s) 104. The machine learning model(s) 104, as described herein, may be trained to compute the boundary points 106 and the class  
25 labels 108 for one or more boundaries.

The method 600, at block B606, includes computing a first set of points that correspond to a traversable boundary. For example, a first set of boundary points may be computed by the machine learning model(s) 104 that correspond to a first boundary dividing drivable free-space from non-drivable space in the real-world environment that may be a  
30 traversable boundary. In some examples, the traversable boundary points may correspond to a single boundary, or may correspond to a portion of the first boundary and a portion of the second boundary.

The method 600, at block B608, includes computing a second set of points that correspond to a non-traversable boundary. For example, a second set of boundary points may be computed by the machine learning model(s) 104 that correspond to a second boundary dividing drivable free-space from non-drivable space in the real-world environment that may be a non-traversable boundary. In some examples, the non-traversable boundary points may correspond to a single boundary, or may correspond to a portion of the first boundary and a portion of the second boundary.

In some examples, the computing the first set of points at block B606 and the computing the second set of points at block B608 may be executed in parallel. For example, one or more GPUs may be used to compute the first set of points in parallel with the second set of points in order to increase processing speeds to enable the machine learning model(s) 104 to run in real-time.

The method 600, at block B610, includes executing one or more operations for controlling the vehicle. For example, one or more operations may be executed for controlling the vehicle 700 within the real-world environment. In some examples, the class labels and real-world locations determined from the locations of the boundary points 106 may be used by the drive stack 146 to determine the one or more operations (e.g., world model updates, trajectories, controls, actuations, etc.). The traversable and non-traversable labels may be useful for the vehicle 700 in determining the one or more operations, such as in emergency situations. For example, in an emergency, with respect to FIG. 5, the vehicle 700 may need to avoid a collision, and to do so, knowing that the boundary segment 502A is traversable as opposed to the non-traversable boundary segment 502B may be useful. In such an example, the vehicle 700 may maneuver over the curb or barrier 516 to avoid collision with a vehicle, pedestrian, or animal, for example, but may not maneuver toward the vehicle 506 in order to avoid a different collision.

The method 600 may repeat, as indicated by the dashed arrow, for each sensor data input.

### **Evaluation Metrics**

In some examples, one or more evaluation metrics may be used to evaluate the machine learning model(s) 104 output and performance. For example, relative gap, precision, label accuracy, smoothness, weighted free-space precision (WFP), weighted free-space recall (WFR), and/or other metrics may be used. Relative gap may measure the average of the absolute deviation between a ground truth boundary curve,  $y_i$  (e.g., the ground

truth boundary points) and a predicted boundary curve,  $\hat{y}_i$  (e.g., the boundary points 106 predicted by the machine learning model(s) 104, by taking a sum of the L1 loss norm ( $\| \cdot \|$ ) and normalizing with the spatial width,  $N$ , of the input image (e.g., the width of the boundary point array) and a spatial height,  $H$ , of the input image. In some non-limiting examples, the

5 relative gap may be determined according to equation (18), below:

$$\text{Relative Gap} = \sum_{i=1}^N \frac{\|y_i - \hat{y}_i\|}{N * H} \quad (18)$$

Precision may measure the amount of overshoot of the predicted boundary curve from the ground truth boundary curve (e.g., the amount of non-drivable space determined to be drivable free-space by the DNN). For example, overshoot should be

10 avoided as non-drivable space may be considered drivable free-space when overshoot is present, which may decrease the safety of the system. As such, precision may be an evaluation metric that is used to ensure that overshoot is either nonexistent or is within acceptable bounds. In some non-limiting examples, precision may be determined according to equation (19), below:

$$\text{Precision} = \frac{\sum_{i=1}^N (H - \max(y_i, \hat{y}_i))}{\sum_{i=1}^N (H - y_i)} \quad (19)$$

where  $H$  is the height of the image,  $y_i$  is the ground truth boundary curve, and  $\hat{y}_i$  is the predicted boundary curve.

Label accuracy may measure the accuracy of the predicted boundary class labels in comparison to the ground truth boundary class labels. In some non-limiting

20 examples, label accuracy may be determined according to equation (20), below:

$$\text{Label Accuracy} = \frac{1}{N_{class}} \sum_{i=1}^{N_{class}} \sum_{j=1}^{N_i} \frac{1(i=\hat{y}_j)}{N_i} \quad (20)$$

where  $N_{class}$  is a total number of boundary classes or types,  $N_i$  is a total number of boundary points labeled as class  $i$ , and  $1(x)$  is an indicator function which is 1 when  $\hat{y}_j$  (class for a boundary point labeled as class  $i$ ) is equal to class  $i$ .

25 Smoothness may measure the smoothness of the predicted curve (e.g., the mean of the difference between consecutive boundary points of the boundary curve). In some non-limiting examples, smoothness may be determined according to equation (21), below:

$$\text{Smoothness} = \sum_{i=1}^{N-1} \frac{\|\hat{y}_{i+1} - \hat{y}_i\|}{N-1} \quad (21)$$

where  $\hat{y}_i$  is the predicted boundary curve,  $\| \cdot \|$  is the L1 norm, and  $N$  is the width of the curve (e.g., the spatial width of the input image).

With respect to WFP and WFR, these evaluation metrics may be based on the concept of drivability (e.g., traversable) and un-drivability (e.g., non-traversable) of obstacles. For example, boundary types or obstacles such as pedestrians or vehicles may have higher un-drivability than a curb. In addition, for WFP and WFR, in some examples, the weight for different boundary class labels may be different, such that some boundary classes or types have a higher associated weight. For example, pedestrians may have a higher associated weight than a curb because accuracy with respect to a human is more pertinent than with respect to a curb.

WFP may include a precision calculation using only overshoot, such that if there is no overshoot, the precision may be 1.0 (e.g., a near perfect match between ground truth and prediction). In some non-limiting examples, WFP may be determined according to equation (22), below:

$$WFP = \frac{(\sum_{p \in \Phi} OFW(p) \cdot F(p) \cdot P(p))}{(\sum_{p \in \Phi} M\_OFW(p) \cdot P(p))} \quad (22)$$

where  $p$  is a boundary point (or pixel),  $\Phi$  is all the boundary points,  $OFW$  is overshoot free-space weight, which may be different based on the boundary class or type,  $F(p)$  is a displacement score for  $p$ ,  $M\_OFW(p)$  is a maximum possible overshoot weight for  $p$ , and  $P(p)$  is in-path weight for  $p$ .

WFR may include a precision calculation using only undershoot, such that if there is no undershoot, the precision may be 1.0 (e.g., a near perfect match between ground truth and prediction). In some non-limiting examples, WFR may be determined according to equation (23), below:

$$WFR = \frac{(\sum_{p \in \Phi} UFW(p) \cdot F(p) \cdot P(p))}{(\sum_{p \in \Phi} M\_UFW(p) \cdot P(p))} \quad (23)$$

where  $p$  is a boundary point (or pixel),  $\Phi$  is all the boundary points,  $UFW$  is undershoot free-space weight, which may be different based on the boundary class or type,  $F(p)$  is a displacement score for  $p$ ,  $M\_UFW(p)$  is a maximum possible undershoot weight for  $p$ , and  $P(p)$  is in-path weight for  $p$ .

### Example Autonomous Vehicle

FIG. 7A is an illustration of an example autonomous vehicle 700, in accordance with some embodiments of the present disclosure. The autonomous vehicle 700

(alternatively referred to herein as the “vehicle 700”) may include a passenger vehicle, such as a car, a truck, a bus, and/or another type of vehicle that accommodates one or more passengers. Autonomous vehicles are generally described in terms of automation levels, defined by the National Highway Traffic Safety Administration (NHTSA), a division of the  
5 US Department of Transportation, and the Society of Automotive Engineers (SAE) “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles” (Standard No. J3016-201806, published on June 15, 2018, Standard No. J3016-201609, published on September 30, 2016, and previous and future versions of this standard). The vehicle 700 may be capable of functionality in accordance with one or more  
10 of Level 3 – Level 5 of the autonomous driving levels. For example, the vehicle 700 may be capable of conditional automation (Level 3), high automation (Level 4), and/or full automation (Level 5), depending on the embodiment.

The vehicle 700 may include components such as a chassis, a vehicle body, wheels (e.g., 2, 4, 6, 8, 18, etc.), tires, axles, and other components of a vehicle. The vehicle  
15 700 may include a propulsion system 750, such as an internal combustion engine, hybrid electric power plant, an all-electric engine, and/or another propulsion system type. The propulsion system 750 may be connected to a drive train of the vehicle 700, which may include a transmission, to enable the propulsion of the vehicle 700. The propulsion system 750 may be controlled in response to receiving signals from the throttle/accelerator 752.

20 A steering system 754, which may include a steering wheel, may be used to steer the vehicle 700 (e.g., along a desired path or route) when the propulsion system 750 is operating (e.g., when the vehicle is in motion). The steering system 754 may receive signals from a steering actuator 756. The steering wheel may be optional for full automation (Level 5) functionality.

25 The brake sensor system 746 may be used to operate the vehicle brakes in response to receiving signals from the brake actuators 748 and/or brake sensors.

Controller(s) 736, which may include one or more system on chips (SoCs) 704 (FIG. 7C) and/or GPU(s), may provide signals (e.g., representative of commands) to one or more components and/or systems of the vehicle 700. For example, the controller(s) may send  
30 signals to operate the vehicle brakes via one or more brake actuators 748, to operate the steering system 754 via one or more steering actuators 756, to operate the propulsion system 750 via one or more throttle/accelerators 752. The controller(s) 736 may include one or more onboard (e.g., integrated) computing devices (e.g., supercomputers) that process sensor

signals, and output operation commands (e.g., signals representing commands) to enable autonomous driving and/or to assist a human driver in driving the vehicle 700. The controller(s) 736 may include a first controller 736 for autonomous driving functions, a second controller 736 for functional safety functions, a third controller 736 for artificial intelligence functionality (e.g., computer vision), a fourth controller 736 for infotainment functionality, a fifth controller 736 for redundancy in emergency conditions, and/or other controllers. In some examples, a single controller 736 may handle two or more of the above functionalities, two or more controllers 736 may handle a single functionality, and/or any combination thereof.

The controller(s) 736 may provide the signals for controlling one or more components and/or systems of the vehicle 700 in response to sensor data received from one or more sensors (e.g., sensor inputs). The sensor data may be received from, for example and without limitation, global navigation satellite systems sensor(s) 758 (e.g., Global Positioning System sensor(s)), RADAR sensor(s) 760, ultrasonic sensor(s) 762, LIDAR sensor(s) 764, inertial measurement unit (IMU) sensor(s) 766 (e.g., accelerometer(s), gyroscope(s), magnetic compass(es), magnetometer(s), etc.), microphone(s) 796, stereo camera(s) 768, wide-view camera(s) 770 (e.g., fisheye cameras), infrared camera(s) 772, surround camera(s) 774 (e.g., 360 degree cameras), long-range and/or mid-range camera(s) 798, speed sensor(s) 744 (e.g., for measuring the speed of the vehicle 700), vibration sensor(s) 742, steering sensor(s) 740, brake sensor(s) (e.g., as part of the brake sensor system 746), and/or other sensor types.

One or more of the controller(s) 736 may receive inputs (e.g., represented by input data) from an instrument cluster 732 of the vehicle 700 and provide outputs (e.g., represented by output data, display data, etc.) via a human-machine interface (HMI) display 734, an audible annunciator, a loudspeaker, and/or via other components of the vehicle 700. The outputs may include information such as vehicle velocity, speed, time, map data (e.g., the HD map 722 of FIG. 7C), location data (e.g., the vehicle's 700 location, such as on a map), direction, location of other vehicles (e.g., an occupancy grid), information about objects and status of objects as perceived by the controller(s) 736, etc. For example, the HMI display 734 may display information about the presence of one or more objects (e.g., a street sign, caution sign, traffic light changing, etc.), and/or information about driving maneuvers the vehicle has made, is making, or will make (e.g., changing lanes now, taking exit 34B in two miles, etc.).

The vehicle 700 further includes a network interface 724 which may use one or more wireless antenna(s) 726 and/or modem(s) to communicate over one or more networks. For example, the network interface 724 may be capable of communication over LTE, WCDMA, UMTS, GSM, CDMA2000, etc. The wireless antenna(s) 726 may also  
5 enable communication between objects in the environment (e.g., vehicles, mobile devices, etc.), using local area network(s), such as Bluetooth, Bluetooth LE, Z-Wave, ZigBee, etc., and/or low power wide-area network(s) (LPWANs), such as LoRaWAN, SigFox, etc.

FIG. 7B is an example of camera locations and fields of view for the example autonomous vehicle 700 of FIG. 7A, in accordance with some embodiments of the present  
10 disclosure. The cameras and respective fields of view are one example embodiment and are not intended to be limiting. For example, additional and/or alternative cameras may be included and/or the cameras may be located at different locations on the vehicle 700.

The camera types for the cameras may include, but are not limited to, digital cameras that may be adapted for use with the components and/or systems of the vehicle 700.  
15 The camera(s) may operate at automotive safety integrity level (ASIL) B and/or at another ASIL. The camera types may be capable of any image capture rate, such as 60 frames per second (fps), 720 fps, 240 fps, etc., depending on the embodiment. The cameras may be capable of using rolling shutters, global shutters, another type of shutter, or a combination thereof. In some examples, the color filter array may include a red clear clear clear (RCCC)  
20 color filter array, a red clear clear blue (RCCB) color filter array, a red blue green clear (RBGC) color filter array, a Foveon X3 color filter array, a Bayer sensors (RGGB) color filter array, a monochrome sensor color filter array, and/or another type of color filter array. In some embodiments, clear pixel cameras, such as cameras with an RCCC, an RCCB, and/or an RBGC color filter array, may be used in an effort to increase light sensitivity.

25 In some examples, one or more of the camera(s) may be used to perform advanced driver assistance systems (ADAS) functions (e.g., as part of a redundant or fail-safe design). For example, a Multi-Function Mono Camera may be installed to provide functions including lane departure warning, traffic sign assist and intelligent headlamp control. One or more of the camera(s) (e.g., all of the cameras) may record and provide image data (e.g.,  
30 video) simultaneously.

One or more of the cameras may be mounted in a mounting assembly, such as a custom designed (3-D printed) assembly, in order to cut out stray light and reflections from within the car (e.g., reflections from the dashboard reflected in the windshield mirrors) which

may interfere with the camera's image data capture abilities. With reference to wing-mirror mounting assemblies, the wing-mirror assemblies may be custom 3-D printed so that the camera mounting plate matches the shape of the wing-mirror. In some examples, the camera(s) may be integrated into the wing-mirror. For side-view cameras, the camera(s) may also be integrated within the four pillars at each corner of the cabin.

Cameras with a field of view that include portions of the environment in front of the vehicle 700 (e.g., front-facing cameras) may be used for surround view, to help identify forward facing paths and obstacles, as well aid in, with the help of one or more controllers 736 and/or control SoCs, providing information critical to generating an occupancy grid and/or determining the preferred vehicle paths. Front-facing cameras may be used to perform many of the same ADAS functions as LIDAR, including emergency braking, pedestrian detection, and collision avoidance. Front-facing cameras may also be used for ADAS functions and systems including Lane Departure Warnings ("LDW"), Autonomous Cruise Control ("ACC"), and/or other functions such as traffic sign recognition.

A variety of cameras may be used in a front-facing configuration, including, for example, a monocular camera platform that includes a CMOS (complementary metal oxide semiconductor) color imager. Another example may be a wide-view camera(s) 770 that may be used to perceive objects coming into view from the periphery (e.g., pedestrians, crossing traffic or bicycles). Although only one wide-view camera is illustrated in FIG. 7B, there may any number of wide-view cameras 770 on the vehicle 700. In addition, long-range camera(s) 798 (e.g., a long-view stereo camera pair) may be used for depth-based object detection, especially for objects for which a neural network has not yet been trained. The long-range camera(s) 798 may also be used for object detection and classification, as well as basic object tracking.

One or more stereo cameras 768 may also be included in a front-facing configuration. The stereo camera(s) 768 may include an integrated control unit comprising a scalable processing unit, which may provide a programmable logic (FPGA) and a multi-core micro-processor with an integrated CAN or Ethernet interface on a single chip. Such a unit may be used to generate a 3-D map of the vehicle's environment, including a distance estimate for all the points in the image. An alternative stereo camera(s) 768 may include a compact stereo vision sensor(s) that may include two camera lenses (one each on the left and right) and an image processing chip that may measure the distance from the vehicle to the target object and use the generated information (e.g., metadata) to activate the autonomous

emergency braking and lane departure warning functions. Other types of stereo camera(s) 768 may be used in addition to, or alternatively from, those described herein.

Cameras with a field of view that include portions of the environment to the side of the vehicle 700 (e.g., side-view cameras) may be used for surround view, providing  
5 information used to create and update the occupancy grid, as well as to generate side impact collision warnings. For example, surround camera(s) 774 (e.g., four surround cameras 774 as illustrated in FIG. 7B) may be positioned to on the vehicle 700. The surround camera(s) 774 may include wide-view camera(s) 770, fisheye camera(s), 360 degree camera(s), and/or the like. For example, four fisheye cameras may be positioned on the vehicle's front, rear, and  
10 sides. In an alternative arrangement, the vehicle may use three surround camera(s) 774 (e.g., left, right, and rear), and may leverage one or more other camera(s) (e.g., a forward-facing camera) as a fourth surround view camera.

Cameras with a field of view that include portions of the environment to the rear of the vehicle 700 (e.g., rear-view cameras) may be used for park assistance, surround  
15 view, rear collision warnings, and creating and updating the occupancy grid. A wide variety of cameras may be used including, but not limited to, cameras that are also suitable as a front-facing camera(s) (e.g., long-range and/or mid-range camera(s) 798, stereo camera(s) 768), infrared camera(s) 772, etc.), as described herein.

FIG. 7C is a block diagram of an example system architecture for the example  
20 autonomous vehicle 700 of FIG. 7A, in accordance with some embodiments of the present disclosure. It should be understood that this and other arrangements described herein are set forth only as examples. Other arrangements and elements (e.g., machines, interfaces, functions, orders, groupings of functions, etc.) may be used in addition to or instead of those shown, and some elements may be omitted altogether. Further, many of the elements  
25 described herein are functional entities that may be implemented as discrete or distributed components or in conjunction with other components, and in any suitable combination and location. Various functions described herein as being performed by entities may be carried out by hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory.

30 Each of the components, features, and systems of the vehicle 700 in FIG. 7C are illustrated as being connected via bus 702. The bus 702 may include a Controller Area Network (CAN) data interface (alternatively referred to herein as a "CAN bus"). A CAN may be a network inside the vehicle 700 used to aid in control of various features and

functionality of the vehicle 700, such as actuation of brakes, acceleration, braking, steering, windshield wipers, etc. A CAN bus may be configured to have dozens or even hundreds of nodes, each with its own unique identifier (e.g., a CAN ID). The CAN bus may be read to find steering wheel angle, ground speed, engine revolutions per minute (RPMs), button  
5 positions, and/or other vehicle status indicators. The CAN bus may be ASIL B compliant.

Although the bus 702 is described herein as being a CAN bus, this is not intended to be limiting. For example, in addition to, or alternatively from, the CAN bus, FlexRay and/or Ethernet may be used. Additionally, although a single line is used to represent the bus 702, this is not intended to be limiting. For example, there may be any  
10 number of busses 702, which may include one or more CAN busses, one or more FlexRay busses, one or more Ethernet busses, and/or one or more other types of busses using a different protocol. In some examples, two or more busses 702 may be used to perform different functions, and/or may be used for redundancy. For example, a first bus 702 may be used for collision avoidance functionality and a second bus 702 may be used for actuation  
15 control. In any example, each bus 702 may communicate with any of the components of the vehicle 700, and two or more busses 702 may communicate with the same components. In some examples, each SoC 704, each controller 736, and/or each computer within the vehicle may have access to the same input data (e.g., inputs from sensors of the vehicle 700), and may be connected to a common bus, such the CAN bus.

20 The vehicle 700 may include one or more controller(s) 736, such as those described herein with respect to FIG. 7A. The controller(s) 736 may be used for a variety of functions. The controller(s) 736 may be coupled to any of the various other components and systems of the vehicle 700, and may be used for control of the vehicle 700, artificial intelligence of the vehicle 700, infotainment for the vehicle 700, and/or the like.

25 The vehicle 700 may include a system(s) on a chip (SoC) 704. The SoC 704 may include CPU(s) 706, GPU(s) 708, processor(s) 710, cache(s) 712, accelerator(s) 714, data store(s) 716, and/or other components and features not illustrated. The SoC(s) 704 may be used to control the vehicle 700 in a variety of platforms and systems. For example, the SoC(s) 704 may be combined in a system (e.g., the system of the vehicle 700) with an HD  
30 map 722 which may obtain map refreshes and/or updates via a network interface 724 from one or more servers (e.g., server(s) 778 of FIG. 7D).

The CPU(s) 706 may include a CPU cluster or CPU complex (alternatively referred to herein as a “CCPLEX”). The CPU(s) 706 may include multiple cores and/or L2

caches. For example, in some embodiments, the CPU(s) 706 may include eight cores in a coherent multi-processor configuration. In some embodiments, the CPU(s) 706 may include four dual-core clusters where each cluster has a dedicated L2 cache (e.g., a 2 MB L2 cache). The CPU(s) 706 (e.g., the CCPLEX) may be configured to support simultaneous cluster  
5 operation enabling any combination of the clusters of the CPU(s) 706 to be active at any given time.

The CPU(s) 706 may implement power management capabilities that include one or more of the following features: individual hardware blocks may be clock-gated automatically when idle to save dynamic power; each core clock may be gated when the core  
10 is not actively executing instructions due to execution of WFI/WFE instructions; each core may be independently power-gated; each core cluster may be independently clock-gated when all cores are clock-gated or power-gated; and/or each core cluster may be independently power-gated when all cores are power-gated. The CPU(s) 706 may further implement an enhanced algorithm for managing power states, where allowed power states  
15 and expected wakeup times are specified, and the hardware/microcode determines the best power state to enter for the core, cluster, and CCPLEX. The processing cores may support simplified power state entry sequences in software with the work offloaded to microcode.

The GPU(s) 708 may include an integrated GPU (alternatively referred to herein as an “iGPU”). The GPU(s) 708 may be programmable and may be efficient for  
20 parallel workloads. The GPU(s) 708, in some examples, may use an enhanced tensor instruction set. The GPU(s) 708 may include one or more streaming microprocessors, where each streaming microprocessor may include an L1 cache (e.g., an L1 cache with at least 96KB storage capacity), and two or more of the streaming microprocessors may share an L2 cache (e.g., an L2 cache with a 512 KB storage capacity). In some embodiments, the GPU(s)  
25 708 may include at least eight streaming microprocessors. The GPU(s) 708 may use compute application programming interface(s) (API(s)). In addition, the GPU(s) 708 may use one or more parallel computing platforms and/or programming models (e.g., NVIDIA’s CUDA).

The GPU(s) 708 may be power-optimized for best performance in automotive and embedded use cases. For example, the GPU(s) 708 may be fabricated on a Fin field-effect transistor (FinFET). However, this is not intended to be limiting and the GPU(s) 708  
30 may be fabricated using other semiconductor manufacturing processes. Each streaming microprocessor may incorporate a number of mixed-precision processing cores partitioned into multiple blocks. For example, and without limitation, 64 PF32 cores and 32 PF64 cores

may be partitioned into four processing blocks. In such an example, each processing block may be allocated 16 FP32 cores, 8 FP64 cores, 16 INT32 cores, two mixed-precision NVIDIA TENSOR COREs for deep learning matrix arithmetic, an L0 instruction cache, a warp scheduler, a dispatch unit, and/or a 64 KB register file. In addition, the streaming microprocessors may include independent parallel integer and floating-point data paths to provide for efficient execution of workloads with a mix of computation and addressing calculations. The streaming microprocessors may include independent thread scheduling capability to enable finer-grain synchronization and cooperation between parallel threads. The streaming microprocessors may include a combined L1 data cache and shared memory unit in order to improve performance while simplifying programming.

The GPU(s) 708 may include a high bandwidth memory (HBM) and/or a 16 GB HBM2 memory subsystem to provide, in some examples, about 900 GB/second peak memory bandwidth. In some examples, in addition to, or alternatively from, the HBM memory, a synchronous graphics random-access memory (SGRAM) may be used, such as a graphics double data rate type five synchronous random-access memory (GDDR5).

The GPU(s) 708 may include unified memory technology including access counters to allow for more accurate migration of memory pages to the processor that accesses them most frequently, thereby improving efficiency for memory ranges shared between processors. In some examples, address translation services (ATS) support may be used to allow the GPU(s) 708 to access the CPU(s) 706 page tables directly. In such examples, when the GPU(s) 708 memory management unit (MMU) experiences a miss, an address translation request may be transmitted to the CPU(s) 706. In response, the CPU(s) 706 may look in its page tables for the virtual-to-physical mapping for the address and transmits the translation back to the GPU(s) 708. As such, unified memory technology may allow a single unified virtual address space for memory of both the CPU(s) 706 and the GPU(s) 708, thereby simplifying the GPU(s) 708 programming and porting of applications to the GPU(s) 708.

In addition, the GPU(s) 708 may include an access counter that may keep track of the frequency of access of the GPU(s) 708 to memory of other processors. The access counter may help ensure that memory pages are moved to the physical memory of the processor that is accessing the pages most frequently.

The SoC(s) 704 may include any number of cache(s) 712, including those described herein. For example, the cache(s) 712 may include an L3 cache that is available to both the CPU(s) 706 and the GPU(s) 708 (e.g., that is connected both the CPU(s) 706 and the

GPU(s) 708). The cache(s) 712 may include a write-back cache that may keep track of states of lines, such as by using a cache coherence protocol (e.g., MEI, MESI, MSI, etc.). The L3 cache may include 4 MB or more, depending on the embodiment, although smaller cache sizes may be used.

5           The SoC(s) 704 may include one or more accelerators 714 (e.g., hardware accelerators, software accelerators, or a combination thereof). For example, the SoC(s) 704 may include a hardware acceleration cluster that may include optimized hardware accelerators and/or large on-chip memory. The large on-chip memory (e.g., 4MB of SRAM), may enable the hardware acceleration cluster to accelerate neural networks and other  
10 calculations. The hardware acceleration cluster may be used to complement the GPU(s) 708 and to off-load some of the tasks of the GPU(s) 708 (e.g., to free up more cycles of the GPU(s) 708 for performing other tasks). As an example, the accelerator(s) 714 may be used for targeted workloads (e.g., perception, convolutional neural networks (CNNs), etc.) that are stable enough to be amenable to acceleration. The term “CNN,” as used herein, may include  
15 all types of CNNs, including region-based or regional convolutional neural networks (RCNNs) and Fast RCNNs (e.g., as used for object detection).

          The accelerator(s) 714 (e.g., the hardware acceleration cluster) may include a deep learning accelerator(s) (DLA). The DLA(s) may include one or more Tensor processing units (TPUs) that may be configured to provide an additional ten trillion operations per  
20 second for deep learning applications and inferencing. The TPUs may be accelerators configured to, and optimized for, performing image processing functions (e.g., for CNNs, RCNNs, etc.). The DLA(s) may further be optimized for a specific set of neural network types and floating point operations, as well as inferencing. The design of the DLA(s) may provide more performance per millimeter than a general-purpose GPU, and vastly exceeds  
25 the performance of a CPU. The TPU(s) may perform several functions, including a single-instance convolution function, supporting, for example, INT8, INT16, and FP16 data types for both features and weights, as well as post-processor functions.

          The DLA(s) may quickly and efficiently execute neural networks, especially CNNs, on processed or unprocessed data for any of a variety of functions, including, for  
30 example and without limitation: a CNN for object identification and detection using data from camera sensors; a CNN for distance estimation using data from camera sensors; a CNN for emergency vehicle detection and identification and detection using data from

microphones; a CNN for facial recognition and vehicle owner identification using data from camera sensors; and/or a CNN for security and/or safety related events.

The DLA(s) may perform any function of the GPU(s) 708, and by using an inference accelerator, for example, a designer may target either the DLA(s) or the GPU(s)  
5 708 for any function. For example, the designer may focus processing of CNNs and floating point operations on the DLA(s) and leave other functions to the GPU(s) 708 and/or other accelerator(s) 714.

The accelerator(s) 714 (e.g., the hardware acceleration cluster) may include a programmable vision accelerator(s) (PVA), which may alternatively be referred to herein as a  
10 computer vision accelerator. The PVA(s) may be designed and configured to accelerate computer vision algorithms for the advanced driver assistance systems (ADAS), autonomous driving, and/or augmented reality (AR) and/or virtual reality (VR) applications. The PVA(s) may provide a balance between performance and flexibility. For example, each PVA(s) may include, for example and without limitation, any number of reduced instruction set computer  
15 (RISC) cores, direct memory access (DMA), and/or any number of vector processors.

The RISC cores may interact with image sensors (e.g., the image sensors of any of the cameras described herein), image signal processor(s), and/or the like. Each of the RISC cores may include any amount of memory. The RISC cores may use any of a number of protocols, depending on the embodiment. In some examples, the RISC cores may execute  
20 a real-time operating system (RTOS). The RISC cores may be implemented using one or more integrated circuit devices, application specific integrated circuits (ASICs), and/or memory devices. For example, the RISC cores may include an instruction cache and/or a tightly coupled RAM.

The DMA may enable components of the PVA(s) to access the system  
25 memory independently of the CPU(s) 706. The DMA may support any number of features used to provide optimization to the PVA including, but not limited to, supporting multi-dimensional addressing and/or circular addressing. In some examples, the DMA may support up to six or more dimensions of addressing, which may include block width, block height, block depth, horizontal block stepping, vertical block stepping, and/or depth stepping.

30 The vector processors may be programmable processors that may be designed to efficiently and flexibly execute programming for computer vision algorithms and provide signal processing capabilities. In some examples, the PVA may include a PVA core and two vector processing subsystem partitions. The PVA core may include a processor subsystem,

DMA engine(s) (e.g., two DMA engines), and/or other peripherals. The vector processing subsystem may operate as the primary processing engine of the PVA, and may include a vector processing unit (VPU), an instruction cache, and/or vector memory (e.g., VMEM). A VPU core may include a digital signal processor such as, for example, a single instruction, multiple data (SIMD), very long instruction word (VLIW) digital signal processor. The combination of the SIMD and VLIW may enhance throughput and speed.

Each of the vector processors may include an instruction cache and may be coupled to dedicated memory. As a result, in some examples, each of the vector processors may be configured to execute independently of the other vector processors. In other examples, the vector processors that are included in a particular PVA may be configured to employ data parallelism. For example, in some embodiments, the plurality of vector processors included in a single PVA may execute the same computer vision algorithm, but on different regions of an image. In other examples, the vector processors included in a particular PVA may simultaneously execute different computer vision algorithms, on the same image, or even execute different algorithms on sequential images or portions of an image. Among other things, any number of PVAs may be included in the hardware acceleration cluster and any number of vector processors may be included in each of the PVAs. In addition, the PVA(s) may include additional error correcting code (ECC) memory, to enhance overall system safety.

The accelerator(s) 714 (e.g., the hardware acceleration cluster) may include a computer vision network on-chip and SRAM, for providing a high-bandwidth, low latency SRAM for the accelerator(s) 714. In some examples, the on-chip memory may include at least 4MB SRAM, consisting of, for example and without limitation, eight field-configurable memory blocks, that may be accessible by both the PVA and the DLA. Each pair of memory blocks may include an advanced peripheral bus (APB) interface, configuration circuitry, a controller, and a multiplexer. Any type of memory may be used. The PVA and DLA may access the memory via a backbone that provides the PVA and DLA with high-speed access to memory. The backbone may include a computer vision network on-chip that interconnects the PVA and the DLA to the memory (e.g., using the APB).

The computer vision network on-chip may include an interface that determines, before transmission of any control signal/address/data, that both the PVA and the DLA provide ready and valid signals. Such an interface may provide for separate phases and separate channels for transmitting control signals/addresses/data, as well as burst-type

communications for continuous data transfer. This type of interface may comply with ISO 26262 or IEC 61508 standards, although other standards and protocols may be used.

In some examples, the SoC(s) 704 may include a real-time ray-tracing hardware accelerator, such as described in U.S. Patent Application No. 16/101,232, filed on  
5 August 10, 2018. The real-time ray-tracing hardware accelerator may be used to quickly and efficiently determine the positions and extents of objects (e.g., within a world model), to generate real-time visualization simulations, for RADAR signal interpretation, for sound propagation synthesis and/or analysis, for simulation of SONAR systems, for general wave propagation simulation, for comparison to LIDAR data for purposes of localization and/or  
10 other functions, and/or for other uses.

The accelerator(s) 714 (e.g., the hardware accelerator cluster) have a wide array of uses for autonomous driving. The PVA may be a programmable vision accelerator that may be used for key processing stages in ADAS and autonomous vehicles. The PVA's capabilities are a good match for algorithmic domains needing predictable processing, at low  
15 power and low latency. In other words, the PVA performs well on semi-dense or dense regular computation, even on small data sets, which need predictable run-times with low latency and low power. Thus, in the context of platforms for autonomous vehicles, the PVAs are designed to run classic computer vision algorithms, as they are efficient at object detection and operating on integer math.

For example, according to one embodiment of the technology, the PVA is  
20 used to perform computer stereo vision. A semi-global matching-based algorithm may be used in some examples, although this is not intended to be limiting. Many applications for Level 3-5 autonomous driving require motion estimation/stereo matching on-the-fly (e.g., structure from motion, pedestrian recognition, lane detection, etc.). The PVA may perform  
25 computer stereo vision function on inputs from two monocular cameras.

In some examples, the PVA may be used to perform dense optical flow. According to process raw RADAR data (e.g., using a 4D Fast Fourier Transform) to provide Processed RADAR. In other examples, the PVA is used for time of flight depth processing, by processing raw time of flight data to provide processed time of flight data, for example.

The DLA may be used to run any type of network to enhance control and  
30 driving safety, including for example, a neural network that outputs a measure of confidence for each object detection. Such a confidence value may be interpreted as a probability, or as providing a relative "weight" of each detection compared to other detections. This

confidence value enables the system to make further decisions regarding which detections should be considered as true positive detections rather than false positive detections. For example, the system may set a threshold value for the confidence and consider only the detections exceeding the threshold value as true positive detections. In an automatic emergency braking (AEB) system, false positive detections would cause the vehicle to automatically perform emergency braking, which is obviously undesirable. Therefore, only the most confident detections should be considered as triggers for AEB. The DLA may run a neural network for regressing the confidence value. The neural network may take as its input at least some subset of parameters, such as bounding box dimensions, ground plane estimate obtained (e.g. from another subsystem), inertial measurement unit (IMU) sensor 766 output that correlates with the vehicle 700 orientation, distance, 3D location estimates of the object obtained from the neural network and/or other sensors (e.g., LIDAR sensor(s) 764 or RADAR sensor(s) 760), among others.

The SoC(s) 704 may include data store(s) 716 (e.g., memory). The data store(s) 716 may be on-chip memory of the SoC(s) 704, which may store neural networks to be executed on the GPU and/or the DLA. In some examples, the data store(s) 716 may be large enough in capacity to store multiple instances of neural networks for redundancy and safety. The data store(s) 712 may comprise L2 or L3 cache(s) 712. Reference to the data store(s) 716 may include reference to the memory associated with the PVA, DLA, and/or other accelerator(s) 714, as described herein.

The SoC(s) 704 may include one or more processor(s) 710 (e.g., embedded processors). The processor(s) 710 may include a boot and power management processor that may be a dedicated processor and subsystem to handle boot power and management functions and related security enforcement. The boot and power management processor may be a part of the SoC(s) 704 boot sequence and may provide runtime power management services. The boot power and management processor may provide clock and voltage programming, assistance in system low power state transitions, management of SoC(s) 704 thermals and temperature sensors, and/or management of the SoC(s) 704 power states. Each temperature sensor may be implemented as a ring-oscillator whose output frequency is proportional to temperature, and the SoC(s) 704 may use the ring-oscillators to detect temperatures of the CPU(s) 706, GPU(s) 708, and/or accelerator(s) 714. If temperatures are determined to exceed a threshold, the boot and power management processor may enter a

temperature fault routine and put the SoC(s) 704 into a lower power state and/or put the vehicle 700 into a chauffeur to safe stop mode (e.g., bring the vehicle 700 to a safe stop).

The processor(s) 710 may further include a set of embedded processors that may serve as an audio processing engine. The audio processing engine may be an audio  
5 subsystem that enables full hardware support for multi-channel audio over multiple interfaces, and a broad and flexible range of audio I/O interfaces. In some examples, the audio processing engine is a dedicated processor core with a digital signal processor with dedicated RAM.

The processor(s) 710 may further include an always on processor engine that  
10 may provide necessary hardware features to support low power sensor management and wake use cases. The always on processor engine may include a processor core, a tightly coupled RAM, supporting peripherals (e.g., timers and interrupt controllers), various I/O controller peripherals, and routing logic.

The processor(s) 710 may further include a safety cluster engine that includes  
15 a dedicated processor subsystem to handle safety management for automotive applications. The safety cluster engine may include two or more processor cores, a tightly coupled RAM, support peripherals (e.g., timers, an interrupt controller, etc.), and/or routing logic. In a safety mode, the two or more cores may operate in a lockstep mode and function as a single core with comparison logic to detect any differences between their operations.

The processor(s) 710 may further include a real-time camera engine that may  
20 include a dedicated processor subsystem for handling real-time camera management.

The processor(s) 710 may further include a high-dynamic range signal processor that may include an image signal processor that is a hardware engine that is part of the camera processing pipeline.

The processor(s) 710 may include a video image compositor that may be a  
25 processing block (e.g., implemented on a microprocessor) that implements video post-processing functions needed by a video playback application to produce the final image for the player window. The video image compositor may perform lens distortion correction on wide-view camera(s) 770, surround camera(s) 774, and/or on in-cabin monitoring camera  
30 sensors. In-cabin monitoring camera sensor is preferably monitored by a neural network running on another instance of the Advanced SoC, configured to identify in cabin events and respond accordingly. An in-cabin system may perform lip reading to activate cellular service and place a phone call, dictate emails, change the vehicle's destination, activate or change the

vehicle's infotainment system and settings, or provide voice-activated web surfing. Certain functions are available to the driver only when the vehicle is operating in an autonomous mode, and are disabled otherwise.

5 The video image compositor may include enhanced temporal noise reduction for both spatial and temporal noise reduction. For example, where motion occurs in a video, the noise reduction weights spatial information appropriately, decreasing the weight of information provided by adjacent frames. Where an image or portion of an image does not include motion, the temporal noise reduction performed by the video image compositor may use information from the previous image to reduce noise in the current image.

10 The video image compositor may also be configured to perform stereo rectification on input stereo lens frames. The video image compositor may further be used for user interface composition when the operating system desktop is in use, and the GPU(s) 708 is not required to continuously render new surfaces. Even when the GPU(s) 708 is powered on and active doing 3D rendering, the video image compositor may be used to  
15 offload the GPU(s) 708 to improve performance and responsiveness.

The SoC(s) 704 may further include a mobile industry processor interface (MIPI) camera serial interface for receiving video and input from cameras, a high-speed interface, and/or a video input block that may be used for camera and related pixel input functions. The SoC(s) 704 may further include an input/output controller(s) that may be  
20 controlled by software and may be used for receiving I/O signals that are uncommitted to a specific role.

The SoC(s) 704 may further include a broad range of peripheral interfaces to enable communication with peripherals, audio codecs, power management, and/or other devices. The SoC(s) 704 may be used to process data from cameras (e.g., connected over  
25 Gigabit Multimedia Serial Link and Ethernet), sensors (e.g., LIDAR sensor(s) 764, RADAR sensor(s) 760, etc. that may be connected over Ethernet), data from bus 702 (e.g., speed of vehicle 700, steering wheel position, etc.), data from GNSS sensor(s) 758 (e.g., connected over Ethernet or CAN bus). The SoC(s) 704 may further include dedicated high-performance mass storage controllers that may include their own DMA engines, and that may be used to  
30 free the CPU(s) 706 from routine data management tasks.

The SoC(s) 704 may be an end-to-end platform with a flexible architecture that spans automation levels 3-5, thereby providing a comprehensive functional safety architecture that leverages and makes efficient use of computer vision and ADAS techniques

for diversity and redundancy, provides a platform for a flexible, reliable driving software stack, along with deep learning tools. The SoC(s) 704 may be faster, more reliable, and even more energy-efficient and space-efficient than conventional systems. For example, the accelerator(s) 714, when combined with the CPU(s) 706, the GPU(s) 708, and the data  
5 store(s) 716, may provide for a fast, efficient platform for level 3-5 autonomous vehicles.

The technology thus provides capabilities and functionality that cannot be achieved by conventional systems. For example, computer vision algorithms may be executed on CPUs, which may be configured using high-level programming language, such as the C programming language, to execute a wide variety of processing algorithms across a  
10 wide variety of visual data. However, CPUs are oftentimes unable to meet the performance requirements of many computer vision applications, such as those related to execution time and power consumption, for example. In particular, many CPUs are unable to execute complex object detection algorithms in real-time, which is a requirement of in-vehicle ADAS applications, and a requirement for practical Level 3-5 autonomous vehicles.

15 In contrast to conventional systems, by providing a CPU complex, GPU complex, and a hardware acceleration cluster, the technology described herein allows for multiple neural networks to be performed simultaneously and/or sequentially, and for the results to be combined together to enable Level 3-5 autonomous driving functionality. For example, a CNN executing on the DLA or dGPU (e.g., the GPU(s) 720) may include a text  
20 and word recognition, allowing the supercomputer to read and understand traffic signs, including signs for which the neural network has not been specifically trained. The DLA may further include a neural network that is able to identify, interpret, and provides semantic understanding of the sign, and to pass that semantic understanding to the path planning modules running on the CPU Complex.

25 As another example, multiple neural networks may be run simultaneously, as is required for Level 3, 4, or 5 driving. For example, a warning sign consisting of “Caution: flashing lights indicate icy conditions,” along with an electric light, may be independently or collectively interpreted by several neural networks. The sign itself may be identified as a traffic sign by a first deployed neural network (e.g., a neural network that has been trained),  
30 the text “Flashing lights indicate icy conditions” may be interpreted by a second deployed neural network, which informs the vehicle’s path planning software (preferably executing on the CPU Complex) that when flashing lights are detected, icy conditions exist. The flashing light may be identified by operating a third deployed neural network over multiple frames,

informing the vehicle's path-planning software of the presence (or absence) of flashing lights. All three neural networks may run simultaneously, such as within the DLA and/or on the GPU(s) 708.

5 In some examples, a CNN for facial recognition and vehicle owner identification may use data from camera sensors to identify the presence of an authorized driver and/or owner of the vehicle 700. The always on sensor processing engine may be used to unlock the vehicle when the owner approaches the driver door and turn on the lights, and, in security mode, to disable the vehicle when the owner leaves the vehicle. In this way, the SoC(s) 704 provide for security against theft and/or carjacking.

10 In another example, a CNN for emergency vehicle detection and identification may use data from microphones 796 to detect and identify emergency vehicle sirens. In contrast to conventional systems, that use general classifiers to detect sirens and manually extract features, the SoC(s) 704 use the CNN for classifying environmental and urban sounds, as well as classifying visual data. In a preferred embodiment, the CNN running on the DLA  
15 is trained to identify the relative closing speed of the emergency vehicle (e.g., by using the Doppler effect). The CNN may also be trained to identify emergency vehicles specific to the local area in which the vehicle is operating, as identified by GNSS sensor(s) 758. Thus, for example, when operating in Europe the CNN will seek to detect European sirens, and when in the United States the CNN will seek to identify only North American sirens. Once an  
20 emergency vehicle is detected, a control program may be used to execute an emergency vehicle safety routine, slowing the vehicle, pulling over to the side of the road, parking the vehicle, and/or idling the vehicle, with the assistance of ultrasonic sensors 762, until the emergency vehicle(s) passes.

The vehicle may include a CPU(s) 718 (e.g., discrete CPU(s), or dCPU(s)),  
25 that may be coupled to the SoC(s) 704 via a high-speed interconnect (e.g., PCIe). The CPU(s) 718 may include an X86 processor, for example. The CPU(s) 718 may be used to perform any of a variety of functions, including arbitrating potentially inconsistent results between ADAS sensors and the SoC(s) 704, and/or monitoring the status and health of the controller(s) 736 and/or infotainment SoC 730, for example.

30 The vehicle 700 may include a GPU(s) 720 (e.g., discrete GPU(s), or dGPU(s)), that may be coupled to the SoC(s) 704 via a high-speed interconnect (e.g., NVIDIA's NVLINK). The GPU(s) 720 may provide additional artificial intelligence functionality, such as by executing redundant and/or different neural networks, and may be

used to train and/or update neural networks based on input (e.g., sensor data) from sensors of the vehicle 700.

The vehicle 700 may further include the network interface 724 which may include one or more wireless antennas 726 (e.g., one or more wireless antennas for different communication protocols, such as a cellular antenna, a Bluetooth antenna, etc.). The network interface 724 may be used to enable wireless connectivity over the Internet with the cloud (e.g., with the server(s) 778 and/or other network devices), with other vehicles, and/or with computing devices (e.g., client devices of passengers). To communicate with other vehicles, a direct link may be established between the two vehicles and/or an indirect link may be established (e.g., across networks and over the Internet). Direct links may be provided using a vehicle-to-vehicle communication link. The vehicle-to-vehicle communication link may provide the vehicle 700 information about vehicles in proximity to the vehicle 700 (e.g., vehicles in front of, on the side of, and/or behind the vehicle 700). This functionality may be part of a cooperative adaptive cruise control functionality of the vehicle 700.

The network interface 724 may include a SoC that provides modulation and demodulation functionality and enables the controller(s) 736 to communicate over wireless networks. The network interface 724 may include a radio frequency front-end for up-conversion from baseband to radio frequency, and down conversion from radio frequency to baseband. The frequency conversions may be performed through well-known processes, and/or may be performed using super-heterodyne processes. In some examples, the radio frequency front end functionality may be provided by a separate chip. The network interface may include wireless functionality for communicating over LTE, WCDMA, UMTS, GSM, CDMA2000, Bluetooth, Bluetooth LE, Wi-Fi, Z-Wave, ZigBee, LoRaWAN, and/or other wireless protocols.

The vehicle 700 may further include data store(s) 728 which may include off-chip (e.g., off the SoC(s) 704) storage. The data store(s) 728 may include one or more storage elements including RAM, SRAM, DRAM, VRAM, Flash, hard disks, and/or other components and/or devices that may store at least one bit of data.

The vehicle 700 may further include GNSS sensor(s) 758. The GNSS sensor(s) 758 (e.g., GPS and/or assisted GPS sensors), to assist in mapping, perception, occupancy grid generation, and/or path planning functions. Any number of GNSS sensor(s) 758 may be used, including, for example and without limitation, a GPS using a USB connector with an Ethernet to Serial (RS-232) bridge.

The vehicle 700 may further include RADAR sensor(s) 760. The RADAR sensor(s) 760 may be used by the vehicle 700 for long-range vehicle detection, even in darkness and/or severe weather conditions. RADAR functional safety levels may be ASIL B. The RADAR sensor(s) 760 may use the CAN and/or the bus 702 (e.g., to transmit data  
5 generated by the RADAR sensor(s) 760) for control and to access object tracking data, with access to Ethernet to access raw data in some examples. A wide variety of RADAR sensor types may be used. For example, and without limitation, the RADAR sensor(s) 760 may be suitable for front, rear, and side RADAR use. In some example, Pulse Doppler RADAR sensor(s) are used.

10 The RADAR sensor(s) 760 may include different configurations, such as long range with narrow field of view, short range with wide field of view, short range side coverage, etc. In some examples, long-range RADAR may be used for adaptive cruise control functionality. The long-range RADAR systems may provide a broad field of view realized by two or more independent scans, such as within a 250m range. The RADAR  
15 sensor(s) 760 may help in distinguishing between static and moving objects, and may be used by ADAS systems for emergency brake assist and forward collision warning. Long-range RADAR sensors may include monostatic multimodal RADAR with multiple (e.g., six or more) fixed RADAR antennae and a high-speed CAN and FlexRay interface. In an example with six antennae, the central four antennae may create a focused beam pattern, designed to  
20 record the vehicle's 700 surroundings at higher speeds with minimal interference from traffic in adjacent lanes. The other two antennae may expand the field of view, making it possible to quickly detect vehicles entering or leaving the vehicle's 700 lane.

Mid-range RADAR systems may include, as an example, a range of up to 760m (front) or 80m (rear), and a field of view of up to 42 degrees (front) or 750 degrees  
25 (rear). Short-range RADAR systems may include, without limitation, RADAR sensors designed to be installed at both ends of the rear bumper. When installed at both ends of the rear bumper, such a RADAR sensor systems may create two beams that constantly monitor the blind spot in the rear and next to the vehicle.

Short-range RADAR systems may be used in an ADAS system for blind spot  
30 detection and/or lane change assist.

The vehicle 700 may further include ultrasonic sensor(s) 762. The ultrasonic sensor(s) 762, which may be positioned at the front, back, and/or the sides of the vehicle 700, may be used for park assist and/or to create and update an occupancy grid. A wide variety of

ultrasonic sensor(s) 762 may be used, and different ultrasonic sensor(s) 762 may be used for different ranges of detection (e.g., 2.5m, 4m). The ultrasonic sensor(s) 762 may operate at functional safety levels of ASIL B.

The vehicle 700 may include LIDAR sensor(s) 764. The LIDAR sensor(s)  
5 764 may be used for object and pedestrian detection, emergency braking, collision avoidance, and/or other functions. The LIDAR sensor(s) 764 may be functional safety level ASIL B. In some examples, the vehicle 700 may include multiple LIDAR sensors 764 (e.g., two, four, six, etc.) that may use Ethernet (e.g., to provide data to a Gigabit Ethernet switch).

In some examples, the LIDAR sensor(s) 764 may be capable of providing a  
10 list of objects and their distances for a 360-degree field of view. Commercially available LIDAR sensor(s) 764 may have an advertised range of approximately 700m, with an accuracy of 2cm-3cm, and with support for a 700Mbps Ethernet connection, for example. In some examples, one or more non-protruding LIDAR sensors 764 may be used. In such examples, the LIDAR sensor(s) 764 may be implemented as a small device that may be  
15 embedded into the front, rear, sides, and/or corners of the vehicle 700. The LIDAR sensor(s) 764, in such examples, may provide up to a 720-degree horizontal and 35-degree vertical field-of-view, with a 200m range even for low-reflectivity objects. Front-mounted LIDAR sensor(s) 764 may be configured for a horizontal field of view between 45 degrees and 135 degrees.

20 In some examples, LIDAR technologies, such as 3D flash LIDAR, may also be used. 3D Flash LIDAR uses a flash of a laser as a transmission source, to illuminate vehicle surroundings up to approximately 200m. A flash LIDAR unit includes a receptor, which records the laser pulse transit time and the reflected light on each pixel, which in turn corresponds to the range from the vehicle to the objects. Flash LIDAR may allow for highly  
25 accurate and distortion-free images of the surroundings to be generated with every laser flash. In some examples, four flash LIDAR sensors may be deployed, one at each side of the vehicle 700. Available 3D flash LIDAR systems include a solid-state 3D staring array LIDAR camera with no moving parts other than a fan (e.g., a non-scanning LIDAR device). The flash LIDAR device may use a 5 nanosecond class I (eye-safe) laser pulse per frame and  
30 may capture the reflected laser light in the form of 3D range point clouds and co-registered intensity data. By using flash LIDAR, and because flash LIDAR is a solid-state device with no moving parts, the LIDAR sensor(s) 764 may be less susceptible to motion blur, vibration, and/or shock.

The vehicle may further include IMU sensor(s) 766. The IMU sensor(s) 766 may be located at a center of the rear axle of the vehicle 700, in some examples. The IMU sensor(s) 766 may include, for example and without limitation, an accelerometer(s), a magnetometer(s), a gyroscope(s), a magnetic compass(es), and/or other sensor types. In  
5 some examples, such as in six-axis applications, the IMU sensor(s) 766 may include accelerometers and gyroscopes, while in nine-axis applications, the IMU sensor(s) 766 may include accelerometers, gyroscopes, and magnetometers.

In some embodiments, the IMU sensor(s) 766 may be implemented as a miniature, high performance GPS-Aided Inertial Navigation System (GPS/INS) that  
10 combines micro-electro-mechanical systems (MEMS) inertial sensors, a high-sensitivity GPS receiver, and advanced Kalman filtering algorithms to provide estimates of position, velocity, and attitude. As such, in some examples, the IMU sensor(s) 766 may enable the vehicle 700 to estimate heading without requiring input from a magnetic sensor by directly observing and correlating the changes in velocity from GPS to the IMU sensor(s) 766. In some examples,  
15 the IMU sensor(s) 766 and the GNSS sensor(s) 758 may be combined in a single integrated unit.

The vehicle may include microphone(s) 796 placed in and/or around the vehicle 700. The microphone(s) 796 may be used for emergency vehicle detection and identification, among other things.

20 The vehicle may further include any number of camera types, including stereo camera(s) 768, wide-view camera(s) 770, infrared camera(s) 772, surround camera(s) 774, long-range and/or mid-range camera(s) 798, and/or other camera types. The cameras may be used to capture image data around an entire periphery of the vehicle 700. The types of cameras used depends on the embodiments and requirements for the vehicle 700, and any  
25 combination of camera types may be used to provide the necessary coverage around the vehicle 700. In addition, the number of cameras may differ depending on the embodiment. For example, the vehicle may include six cameras, seven cameras, ten cameras, twelve cameras, and/or another number of cameras. The cameras may support, as an example and without limitation, Gigabit Multimedia Serial Link (GMSL) and/or Gigabit Ethernet. Each of  
30 the camera(s) is described with more detail herein with respect to FIG. 7A and FIG. 7B.

The vehicle 700 may further include vibration sensor(s) 742. The vibration sensor(s) 742 may measure vibrations of components of the vehicle, such as the axle(s). For example, changes in vibrations may indicate a change in road surfaces. In another example,

when two or more vibration sensors 742 are used, the differences between the vibrations may be used to determine friction or slippage of the road surface (e.g., when the difference in vibration is between a power-driven axle and a freely rotating axle).

The vehicle 700 may include an ADAS system 738. The ADAS system 738  
5 may include a SoC, in some examples. The ADAS system 738 may include autonomous/adaptive/automatic cruise control (ACC), cooperative adaptive cruise control (CACC), forward crash warning (FCW), automatic emergency braking (AEB), lane departure warnings (LDW), lane keep assist (LKA), blind spot warning (BSW), rear cross-traffic warning (RCTW), collision warning systems (CWS), lane centering (LC), and/or other  
10 features and functionality.

The ACC systems may use RADAR sensor(s) 760, LIDAR sensor(s) 764, and/or a camera(s). The ACC systems may include longitudinal ACC and/or lateral ACC. Longitudinal ACC monitors and controls the distance to the vehicle immediately ahead of the vehicle 700 and automatically adjust the vehicle speed to maintain a safe distance from  
15 vehicles ahead. Lateral ACC performs distance keeping, and advises the vehicle 700 to change lanes when necessary. Lateral ACC is related to other ADAS applications such as LCA and CWS.

CACC uses information from other vehicles that may be received via the network interface 724 and/or the wireless antenna(s) 726 from other vehicles via a wireless  
20 link, or indirectly, over a network connection (e.g., over the Internet). Direct links may be provided by a vehicle-to-vehicle (V2V) communication link, while indirect links may be infrastructure-to-vehicle (I2V) communication link. In general, the V2V communication concept provides information about the immediately preceding vehicles (e.g., vehicles immediately ahead of and in the same lane as the vehicle 700), while the I2V communication  
25 concept provides information about traffic further ahead. CACC systems may include either or both I2V and V2V information sources. Given the information of the vehicles ahead of the vehicle 700, CACC may be more reliable and it has potential to improve traffic flow smoothness and reduce congestion on the road.

FCW systems are designed to alert the driver to a hazard, so that the driver  
30 may take corrective action. FCW systems use a front-facing camera and/or RADAR sensor(s) 760, coupled to a dedicated processor, DSP, FPGA, and/or ASIC, that is electrically coupled to driver feedback, such as a display, speaker, and/or vibrating component. FCW

systems may provide a warning, such as in the form of a sound, visual warning, vibration and/or a quick brake pulse.

AEB systems detect an impending forward collision with another vehicle or other object, and may automatically apply the brakes if the driver does not take corrective  
5 action within a specified time or distance parameter. AEB systems may use front-facing camera(s) and/or RADAR sensor(s) 760, coupled to a dedicated processor, DSP, FPGA, and/or ASIC. When the AEB system detects a hazard, it typically first alerts the driver to take corrective action to avoid the collision and, if the driver does not take corrective action, the AEB system may automatically apply the brakes in an effort to prevent, or at least  
10 mitigate, the impact of the predicted collision. AEB systems, may include techniques such as dynamic brake support and/or crash imminent braking.

LDW systems provide visual, audible, and/or tactile warnings, such as steering wheel or seat vibrations, to alert the driver when the vehicle 700 crosses lane markings. A LDW system does not activate when the driver indicates an intentional lane departure, by  
15 activating a turn signal. LDW systems may use front-side facing cameras, coupled to a dedicated processor, DSP, FPGA, and/or ASIC, that is electrically coupled to driver feedback, such as a display, speaker, and/or vibrating component.

LKA systems are a variation of LDW systems. LKA systems provide steering input or braking to correct the vehicle 700 if the vehicle 700 starts to exit the lane.

BSW systems detects and warn the driver of vehicles in an automobile's blind  
20 spot. BSW systems may provide a visual, audible, and/or tactile alert to indicate that merging or changing lanes is unsafe. The system may provide an additional warning when the driver uses a turn signal. BSW systems may use rear-side facing camera(s) and/or RADAR sensor(s) 760, coupled to a dedicated processor, DSP, FPGA, and/or ASIC, that is electrically  
25 coupled to driver feedback, such as a display, speaker, and/or vibrating component.

RCTW systems may provide visual, audible, and/or tactile notification when an object is detected outside the rear-camera range when the vehicle 700 is backing up. Some RCTW systems include AEB to ensure that the vehicle brakes are applied to avoid a crash. RCTW systems may use one or more rear-facing RADAR sensor(s) 760, coupled to a  
30 dedicated processor, DSP, FPGA, and/or ASIC, that is electrically coupled to driver feedback, such as a display, speaker, and/or vibrating component.

Conventional ADAS systems may be prone to false positive results which may be annoying and distracting to a driver, but typically are not catastrophic, because the ADAS

systems alert the driver and allow the driver to decide whether a safety condition truly exists and act accordingly. However, in an autonomous vehicle 700, the vehicle 700 itself must, in the case of conflicting results, decide whether to heed the result from a primary computer or a secondary computer (e.g., a first controller 736 or a second controller 736). For example, in  
5 some embodiments, the ADAS system 738 may be a backup and/or secondary computer for providing perception information to a backup computer rationality module. The backup computer rationality monitor may run a redundant diverse software on hardware components to detect faults in perception and dynamic driving tasks. Outputs from the ADAS system 738 may be provided to a supervisory MCU. If outputs from the primary computer and the  
10 secondary computer conflict, the supervisory MCU must determine how to reconcile the conflict to ensure safe operation.

In some examples, the primary computer may be configured to provide the supervisory MCU with a confidence score, indicating the primary computer's confidence in the chosen result. If the confidence score exceeds a threshold, the supervisory MCU may  
15 follow the primary computer's direction, regardless of whether the secondary computer provides a conflicting or inconsistent result. Where the confidence score does not meet the threshold, and where the primary and secondary computer indicate different results (e.g., the conflict), the supervisory MCU may arbitrate between the computers to determine the appropriate outcome.

20 The supervisory MCU may be configured to run a neural network(s) that is trained and configured to determine, based on outputs from the primary computer and the secondary computer, conditions under which the secondary computer provides false alarms. Thus, the neural network(s) in the supervisory MCU may learn when the secondary computer's output may be trusted, and when it cannot. For example, when the secondary  
25 computer is a RADAR-based FCW system, a neural network(s) in the supervisory MCU may learn when the FCW system is identifying metallic objects that are not, in fact, hazards, such as a drainage grate or manhole cover that triggers an alarm. Similarly, when the secondary computer is a camera-based LDW system, a neural network in the supervisory MCU may learn to override the LDW when bicyclists or pedestrians are present and a lane departure is,  
30 in fact, the safest maneuver. In embodiments that include a neural network(s) running on the supervisory MCU, the supervisory MCU may include at least one of a DLA or GPU suitable for running the neural network(s) with associated memory. In preferred embodiments, the supervisory MCU may comprise and/or be included as a component of the SoC(s) 704.

In other examples, ADAS system 738 may include a secondary computer that performs ADAS functionality using traditional rules of computer vision. As such, the secondary computer may use classic computer vision rules (if-then), and the presence of a neural network(s) in the supervisory MCU may improve reliability, safety and performance.

5 For example, the diverse implementation and intentional non-identity makes the overall system more fault-tolerant, especially to faults caused by software (or software-hardware interface) functionality. For example, if there is a software bug or error in the software running on the primary computer, and the non-identical software code running on the secondary computer provides the same overall result, the supervisory MCU may have greater

10 confidence that the overall result is correct, and the bug in software or hardware on primary computer is not causing material error.

In some examples, the output of the ADAS system 738 may be fed into the primary computer's perception block and/or the primary computer's dynamic driving task block. For example, if the ADAS system 738 indicates a forward crash warning due to an

15 object immediately ahead, the perception block may use this information when identifying objects. In other examples, the secondary computer may have its own neural network which is trained and thus reduces the risk of false positives, as described herein.

The vehicle 700 may further include the infotainment SoC 730 (e.g., an in-vehicle infotainment system (IVI)). Although illustrated and described as a SoC, the

20 infotainment system may not be a SoC, and may include two or more discrete components. The infotainment SoC 730 may include a combination of hardware and software that may be used to provide audio (e.g., music, a personal digital assistant, navigational instructions, news, radio, etc.), video (e.g., TV, movies, streaming, etc.), phone (e.g., hands-free calling), network connectivity (e.g., LTE, Wi-Fi, etc.), and/or information services (e.g., navigation

25 systems, rear-parking assistance, a radio data system, vehicle related information such as fuel level, total distance covered, brake fuel level, oil level, door open/close, air filter information, etc.) to the vehicle 700. For example, the infotainment SoC 730 may radios, disk players, navigation systems, video players, USB and Bluetooth connectivity, carputers, in-car entertainment, Wi-Fi, steering wheel audio controls, hands free voice control, a heads-up

30 display (HUD), an HMI display 734, a telematics device, a control panel (e.g., for controlling and/or interacting with various components, features, and/or systems), and/or other components. The infotainment SoC 730 may further be used to provide information (e.g., visual and/or audible) to a user(s) of the vehicle, such as information from the ADAS system

738, autonomous driving information such as planned vehicle maneuvers, trajectories, surrounding environment information (e.g., intersection information, vehicle information, road information, etc.), and/or other information.

The infotainment SoC 730 may include GPU functionality. The infotainment  
5 SoC 730 may communicate over the bus 702 (e.g., CAN bus, Ethernet, etc.) with other devices, systems, and/or components of the vehicle 700. In some examples, the infotainment SoC 730 may be coupled to a supervisory MCU such that the GPU of the infotainment system may perform some self-driving functions in the event that the primary controller(s) 736 (e.g., the primary and/or backup computers of the vehicle 700) fail. In such an example,  
10 the infotainment SoC 730 may put the vehicle 700 into a chauffeur to safe stop mode, as described herein.

The vehicle 700 may further include an instrument cluster 732 (e.g., a digital dash, an electronic instrument cluster, a digital instrument panel, etc.). The instrument cluster 732 may include a controller and/or supercomputer (e.g., a discrete controller or  
15 supercomputer). The instrument cluster 732 may include a set of instrumentation such as a speedometer, fuel level, oil pressure, tachometer, odometer, turn indicators, gearshift position indicator, seat belt warning light(s), parking-brake warning light(s), engine-malfunction light(s), airbag (SRS) system information, lighting controls, safety system controls, navigation information, etc. In some examples, information may be displayed and/or shared  
20 among the infotainment SoC 730 and the instrument cluster 732. In other words, the instrument cluster 732 may be included as part of the infotainment SoC 730, or vice versa.

FIG. 7D is a system diagram for communication between cloud-based server(s) and the example autonomous vehicle 700 of FIG. 7A, in accordance with some embodiments of the present disclosure. The system 776 may include server(s) 778,  
25 network(s) 790, and vehicles, including the vehicle 700. The server(s) 778 may include a plurality of GPUs 784(A)-784(H) (collectively referred to herein as GPUs 784), PCIe switches 782(A)-782(H) (collectively referred to herein as PCIe switches 782), and/or CPUs 780(A)-780(B) (collectively referred to herein as CPUs 780). The GPUs 784, the CPUs 780, and the PCIe switches may be interconnected with high-speed interconnects such as, for  
30 example and without limitation, NVLink interfaces 788 developed by NVIDIA and/or PCIe connections 786. In some examples, the GPUs 784 are connected via NVLink and/or NVSwitch SoC and the GPUs 784 and the PCIe switches 782 are connected via PCIe interconnects. Although eight GPUs 784, two CPUs 780, and two PCIe switches are

illustrated, this is not intended to be limiting. Depending on the embodiment, each of the server(s) 778 may include any number of GPUs 784, CPUs 780, and/or PCIe switches. For example, the server(s) 778 may each include eight, sixteen, thirty-two, and/or more GPUs 784.

5                   The server(s) 778 may receive, over the network(s) 790 and from the vehicles, image data representative of images showing unexpected or changed road conditions, such as recently commenced road-work. The server(s) 778 may transmit, over the network(s) 790 and to the vehicles, neural networks 792, updated neural networks 792, and/or map information 794, including information regarding traffic and road conditions. The updates to  
10 the map information 794 may include updates for the HD map 722, such as information regarding construction sites, potholes, detours, flooding, and/or other obstructions. In some examples, the neural networks 792, the updated neural networks 792, and/or the map information 794 may have resulted from new training and/or experiences represented in data received from any number of vehicles in the environment, and/or based on training performed  
15 at a datacenter (e.g., using the server(s) 778 and/or other servers).

                  The server(s) 778 may be used to train machine learning models (e.g., neural networks) based on training data. The training data may be generated by the vehicles, and/or may be generated in a simulation (e.g., using a game engine). In some examples, the training data is tagged (e.g., where the neural network benefits from supervised learning) and/or  
20 undergoes other pre-processing, while in other examples the training data is not tagged and/or pre-processed (e.g., where the neural network does not require supervised learning). Once the machine learning models are trained, the machine learning models may be used by the vehicles (e.g., transmitted to the vehicles over the network(s) 790, and/or the machine learning models may be used by the server(s) 778 to remotely monitor the vehicles.

25                   In some examples, the server(s) 778 may receive data from the vehicles and apply the data to up-to-date real-time neural networks for real-time intelligent inferencing. The server(s) 778 may include deep-learning supercomputers and/or dedicated AI computers powered by GPU(s) 784, such as a DGX and DGX Station machines developed by NVIDIA. However, in some examples, the server(s) 778 may include deep learning infrastructure that  
30 use only CPU-powered datacenters.

                  The deep-learning infrastructure of the server(s) 778 may be capable of fast, real-time inferencing, and may use that capability to evaluate and verify the health of the processors, software, and/or associated hardware in the vehicle 700. For example, the deep-

- 70 -

learning infrastructure may receive periodic updates from the vehicle 700, such as a sequence of images and/or objects that the vehicle 700 has located in that sequence of images (e.g., via computer vision and/or other machine learning object classification techniques). The deep-learning infrastructure may run its own neural network to identify the objects and compare  
5 them with the objects identified by the vehicle 700 and, if the results do not match and the infrastructure concludes that the AI in the vehicle 700 is malfunctioning, the server(s) 778 may transmit a signal to the vehicle 700 instructing a fail-safe computer of the vehicle 700 to assume control, notify the passengers, and complete a safe parking maneuver.

For inferencing, the server(s) 778 may include the GPU(s) 784 and one or  
10 more programmable inference accelerators (e.g., NVIDIA's TensorRT 3). The combination of GPU-powered servers and inference acceleration may make real-time responsiveness possible. In other examples, such as where performance is less critical, servers powered by CPUs, FPGAs, and other processors may be used for inferencing.

#### **Example Computing Device**

FIG. 8 is a block diagram of an example computing device 800 suitable for  
15 use in implementing some embodiments of the present disclosure. Computing device 800 may include a bus 802 that directly or indirectly couples the following devices: memory 804, one or more central processing units (CPUs) 806, one or more graphics processing units (GPUs) 808, a communication interface 810, input/output (I/O) ports 812, input/output  
20 components 814, a power supply 816, and one or more presentation components 818 (e.g., display(s)).

Although the various blocks of FIG. 8 are shown as connected via the bus 802 with lines, this is not intended to be limiting and is for clarity only. For example, in some embodiments, a presentation component 818, such as a display device, may be considered an  
25 I/O component 814 (e.g., if the display is a touch screen). As another example, the CPUs 806 and/or GPUs 808 may include memory (e.g., the memory 804 may be representative of a storage device in addition to the memory of the GPUs 808, the CPUs 806, and/or other components). In other words, the computing device of FIG. 8 is merely illustrative. Distinction is not made between such categories as "workstation," "server," "laptop,"  
30 "desktop," "tablet," "client device," "mobile device," "hand-held device," "game console," "electronic control unit (ECU)," "virtual reality system," and/or other device or system types, as all are contemplated within the scope of the computing device of FIG. 8.

The bus 802 may represent one or more busses, such as an address bus, a data bus, a control bus, or a combination thereof. The bus 802 may include one or more bus types, such as an industry standard architecture (ISA) bus, an extended industry standard architecture (EISA) bus, a video electronics standards association (VESA) bus, a peripheral component interconnect (PCI) bus, a peripheral component interconnect express (PCIe) bus, and/or another type of bus.

The memory 804 may include any of a variety of computer-readable media. The computer-readable media may be any available media that may be accessed by the computing device 800. The computer-readable media may include both volatile and nonvolatile media, and removable and non-removable media. By way of example, and not limitation, the computer-readable media may comprise computer-storage media and communication media.

The computer-storage media may include both volatile and nonvolatile media and/or removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules, and/or other data types. For example, the memory 804 may store computer-readable instructions (e.g., that represent a program(s) and/or a program element(s), such as an operating system. Computer-storage media may include, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which may be used to store the desired information and which may be accessed by computing device 800. As used herein, computer storage media does not comprise signals per se.

The communication media may embody computer-readable instructions, data structures, program modules, and/or other data types in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term “modulated data signal” may refer to a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, the communication media may include wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer-readable media.

The CPU(s) 806 may be configured to execute the computer-readable instructions to control one or more components of the computing device 800 to perform one or more of the methods and/or processes described herein. The CPU(s) 806 may each include one or more cores (e.g., one, two, four, eight, twenty-eight, seventy-two, etc.) that are capable of handling a multitude of software threads simultaneously. The CPU(s) 806 may include any type of processor, and may include different types of processors depending on the type of computing device 800 implemented (e.g., processors with fewer cores for mobile devices and processors with more cores for servers). For example, depending on the type of computing device 800, the processor may be an ARM processor implemented using Reduced Instruction Set Computing (RISC) or an x86 processor implemented using Complex Instruction Set Computing (CISC). The computing device 800 may include one or more CPUs 806 in addition to one or more microprocessors or supplementary co-processors, such as math co-processors.

The GPU(s) 808 may be used by the computing device 800 to render graphics (e.g., 3D graphics). The GPU(s) 808 may include hundreds or thousands of cores that are capable of handling hundreds or thousands of software threads simultaneously. The GPU(s) 808 may generate pixel data for output images in response to rendering commands (e.g., rendering commands from the CPU(s) 806 received via a host interface). The GPU(s) 808 may include graphics memory, such as display memory, for storing pixel data. The display memory may be included as part of the memory 804. The GPU(s) 708 may include two or more GPUs operating in parallel (e.g., via a link). When combined together, each GPU 808 may generate pixel data for different portions of an output image or for different output images (e.g., a first GPU for a first image and a second GPU for a second image). Each GPU may include its own memory, or may share memory with other GPUs.

In examples where the computing device 800 does not include the GPU(s) 808, the CPU(s) 806 may be used to render graphics.

The communication interface 810 may include one or more receivers, transmitters, and/or transceivers that enable the computing device 700 to communicate with other computing devices via an electronic communication network, included wired and/or wireless communications. The communication interface 810 may include components and functionality to enable communication over any of a number of different networks, such as wireless networks (e.g., Wi-Fi, Z-Wave, Bluetooth, Bluetooth LE, ZigBee, etc.), wired

networks (e.g., communicating over Ethernet), low-power wide-area networks (e.g., LoRaWAN, SigFox, etc.), and/or the Internet.

The I/O ports 812 may enable the computing device 800 to be logically coupled to other devices including the I/O components 814, the presentation component(s) 818, and/or other components, some of which may be built in to (e.g., integrated in) the computing device 800. Illustrative I/O components 814 include a microphone, mouse, keyboard, joystick, game pad, game controller, satellite dish, scanner, printer, wireless device, etc. The I/O components 814 may provide a natural user interface (NUI) that processes air gestures, voice, or other physiological inputs generated by a user. In some instances, inputs may be transmitted to an appropriate network element for further processing. An NUI may implement any combination of speech recognition, stylus recognition, facial recognition, biometric recognition, gesture recognition both on screen and adjacent to the screen, air gestures, head and eye tracking, and touch recognition (as described in more detail below) associated with a display of the computing device 800. The computing device 800 may include depth cameras, such as stereoscopic camera systems, infrared camera systems, RGB camera systems, touchscreen technology, and combinations of these, for gesture detection and recognition. Additionally, the computing device 800 may include accelerometers or gyroscopes (e.g., as part of an inertia measurement unit (IMU)) that enable detection of motion. In some examples, the output of the accelerometers or gyroscopes may be used by the computing device 800 to render immersive augmented reality or virtual reality.

The power supply 816 may include a hard-wired power supply, a battery power supply, or a combination thereof. The power supply 816 may provide power to the computing device 800 to enable the components of the computing device 800 to operate.

The presentation component(s) 818 may include a display (e.g., a monitor, a touch screen, a television screen, a heads-up-display (HUD), other display types, or a combination thereof), speakers, and/or other presentation components. The presentation component(s) 818 may receive data from other components (e.g., the GPU(s) 808, the CPU(s) 806, etc.), and output the data (e.g., as an image, video, sound, etc.).

The disclosure may be described in the general context of computer code or machine-useable instructions, including computer-executable instructions such as program modules, being executed by a computer or other machine, such as a personal data assistant or other handheld device. Generally, program modules including routines, programs, objects,

components, data structures, etc., refer to code that perform particular tasks or implement particular abstract data types. The disclosure may be practiced in a variety of system configurations, including hand-held devices, consumer electronics, general-purpose computers, more specialty computing devices, etc. The disclosure may also be practiced in  
5 distributed computing environments where tasks are performed by remote-processing devices that are linked through a communications network.

As used herein, a recitation of “and/or” with respect to two or more elements should be interpreted to mean only one element, or a combination of elements. For example, “element A, element B, and/or element C” may include only element A, only element B, only  
10 element C, element A and element B, element A and element C, element B and element C, or elements A, B, and C. In addition, “at least one of element A or element B” may include at least one of element A, at least one of element B, or at least one of element A and at least one of element B. Further, “at least one of element A and element B” may include at least one of  
15 element A, at least one of element B, or at least one of element A and at least one of element B.

The subject matter of the present disclosure is described with specificity herein to meet statutory requirements. However, the description itself is not intended to limit the scope of this disclosure. Rather, the inventors have contemplated that the claimed subject matter might also be embodied in other ways, to include different steps or combinations of  
20 steps similar to the ones described in this document, in conjunction with other present or future technologies. Moreover, although the terms “step” and/or “block” may be used herein to connote different elements of methods employed, the terms should not be interpreted as implying any particular order among or between various steps herein disclosed unless and except when the order of individual steps is explicitly described.

25

## CLAIMS

What is claimed is:

1. A method comprising: receiving sensor data generated by one or more sensors of a vehicle in a physical environment; applying the sensor data to a neural network;  
5 computing, by a first stream of the neural network, a set of boundary points represented by the sensor data that correspond to a boundary dividing drivable free-space within the physical environment from non-drivable space within the physical environment; determining locations within the physical environment corresponding to the set of boundary points; computing, by a second stream of the neural network, class labels for boundary points of the set of boundary  
10 points; and executing one or more operations for controlling the vehicle through the physical environment based at least in part on the locations and the class labels.
2. The method of claim 1, wherein the set of boundary points is a first set of boundary points, the sensor data is first sensor data, and the method further comprises: weighting first values corresponding to the first set of boundary points with second values  
15 corresponding to a second set of boundary points from second sensor data captured prior to the first sensor data to compute third values, wherein the determining the locations within the physical environment is based at least in part on the third values.
3. The method of claim 1, further comprising: for at least a boundary point of the set of boundary points, applying a filter to a first value corresponding to the boundary  
20 point to generate an updated first value, the filter weighting a second value corresponding to at least one adjacent boundary point of the boundary point, wherein the determining the locations within the physical environment is based at least in part on the updated first value.
4. The method of claim 1, wherein the class labels include at least one of pedestrian, curb, barrier, vehicle, structure, traversable, non-traversable, or a catchall.
- 25 5. The method of claim 1, wherein the neural network includes a convolutional neural network (CNN), the CNN performs a regression analysis, and the set of boundary points is computed by the first stream of the CNN as a one-dimensional array.

- 76 -

6. The method of claim 5, wherein the sensor data includes image data representative of an image having a spatial width, and the one-dimensional array includes a number of columns corresponding to the spatial width.

7. The method of claim 1, wherein the set of boundary points is a first set of  
5 boundary points, the drivable free-space is a first drivable free-space, the non-drivable space is a first non-drivable space, and the method further comprises: computing, by a third stream of the neural network, a second set of boundary points represented by the sensor data that corresponds to a second boundary dividing a second drivable free-space within the physical environment from a second non-drivable space within the physical environment.

10 8. The method of claim 1, wherein the sensor data is representative of a number of columns, and the set of boundary points includes a boundary point from each column of the number of columns.

9. The method of claim 8, wherein the boundary point from each column is a single point, and the single point represents, ascending from the bottom of an image  
15 represented by the sensor data, a first point in the column corresponding to the boundary.

10. A method comprising: receiving a training image representative of a physical environment; receiving annotation data representative of one or more annotations corresponding to the training image, the one or more annotations including: a boundary label extending across a width of the training image and corresponding to at least one pixel in each  
20 column of pixels of the training image, the at least one pixel corresponding to, from bottom to top of the training image, a first boundary in the environment at least partially dividing drivable free-space from non-drivable space; and class labels corresponding to one or more segments of the boundary label, each class label corresponding to a boundary type associated with the one or more segments; and training a neural network using the one or more  
25 annotations as ground truth data.

11. The method of claim 10, wherein a segment of the one or more segments of the boundary label corresponds to a shape defined by the boundary type.

- 77 -

12. The method of claim 10, wherein a segment of the one or more segments of the boundary label includes a straight line extending along a portion of the training image that corresponds to the boundary type.

13. The method of claim 10, wherein the training of the neural network  
5 includes assigning a higher weight to output data from the neural network that has a first class label of the class labels than a second class label, and adjusting the neural network using the higher weight.

14. The method of claim 10, wherein: the neural network computes a first output corresponding to the class labels and a second output corresponding to a location of  
10 the boundary label within the training image; and the training the neural network includes using a first loss function associated with the first output and a second loss function associated with the second output.

15. The method of claim 14, wherein the training the neural network includes using a weighted combination of the first loss function and the second loss function to  
15 compute a final loss.

16. A method comprising: receiving sensor data representative of an image of a field of view of at least one sensor of a vehicle in a physical environment; applying the sensor data to a neural network; computing, by the neural network, a first set of boundary points represented by the sensor data that corresponds to a traversable boundary dividing first  
20 drivable free-space within the physical environment from first non-drivable space within the physical environment; computing, by the neural network, a second set of boundary points represented by the sensor data that corresponds to a non-traversable boundary dividing second drivable free-space within the physical environment from at least one of the first non-drivable space or second non-drivable space within the physical environment; and executing  
25 one or more operations for controlling the vehicle through the physical environment based at least in part on the traversable boundary and the non-traversable boundary.

17. The method of claim 16, wherein the controlling the vehicle through the physical environment includes traversing the traversable boundary.

- 78 -

18. The method of claim 16, wherein the traversable boundary corresponds to a first segment of a larger boundary and the non-traversable boundary corresponds to a second segment of the larger boundary.

19. The method of claim 16, wherein the non-traversable boundary is located  
5 further from the vehicle in the physical environment than the traversable boundary.

20. The method of claim 16, wherein the neural network is trained to compute a plurality of boundary classes, the plurality of boundary classes including at least traversable and non-traversable.




FIGURE 1A




FIGURE 1B




FIGURE 1C





FIGURE 1D

5/15

**FIGURE 2**

300A




FIGURE 3A

300B






FIGURE 3B


8/15


**FIGURE 4**



## FIGURE 5

10/15

**FIGURE 6**



# FIGURE 7A




FIGURE 7B

13/15




FIGURE 7C

776 ↗




FIGURE 7D




FIGURE 8

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2019/022592

A. CLASSIFICATION OF SUBJECT MATTER  
INV. G06K9/00 G06K9/66 G06K9/62  
ADD.

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

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
G06K

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

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

EPO-Internal, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                       | Relevant to claim No. |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| X         | GARNETT NOA ET AL: "Real-Time Category-Based and General Obstacle Detection for Autonomous Driving", 2017 IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION WORKSHOPS (ICCVW), IEEE, 22 October 2017 (2017-10-22), pages 198-205, XP033303458, DOI: 10.1109/ICCVW.2017.32 | 1-12                  |
| A         | abstract<br>figures 1,3,4,5<br>section 2 - Combined network for obstacle detection                                                                                                                                                                                       | 13-15                 |
| X         | US 2016/247290 A1 (LIU MING-YU [US] ET AL)<br>25 August 2016 (2016-08-25)<br>figures 1-4<br>paragraph [0017] - paragraph [0036]                                                                                                                                          | 1-15                  |



Further documents are listed in the continuation of Box C.



See patent family annex.

\* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

17 June 2019

Date of mailing of the international search report

26/08/2019

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040,  
Fax: (+31-70) 340-3016

Authorized officer

Borges, Pedro

# INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US2019/022592

## Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

## Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-15

### Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

**FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210**

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-15

Neural network with first stream for computing boundary and second stream for computing class labels.

---

2. claims: 16-20

Calculating two sets of boundary points by neural network.

---

## INTERNATIONAL SEARCH REPORT

### Information on patent family members

International application No

PCT/US2019/022592

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date |
|-------------------------------------------|---------------------|----------------------------|---------------------|
| US 2016247290 A1                          | 25-08-2016          | JP 6485646 B2              | 20-03-2019          |
|                                           |                     | JP 2016157434 A            | 01-09-2016          |
|                                           |                     | US 2016247290 A1           | 25-08-2016          |
| -----                                     |                     |                            |                     |