

ChipNet: Real-Time LiDAR Processing for Drivable Region Segmentation on an FPGA

Yecheng Lyu, Lin Bai and Xinming Huang *Senior Member, IEEE*

Abstract—This paper presents a field-programmable gate array (FPGA) design of a segmentation algorithm based on convolutional neural network (CNN) that can process light detection and ranging (LiDAR) data in real-time. For autonomous vehicles, drivable region segmentation is an essential step that sets up the static constraints for planning tasks. Traditional drivable region segmentation algorithms are mostly developed on camera data, so their performance is susceptible to the light conditions and the qualities of road markings. LiDAR sensors can obtain the 3D geometry information of the vehicle surroundings with high precision. However, it is a computational challenge to process a large amount of LiDAR data in real-time. In this paper, a convolutional neural network model is proposed and trained to perform semantic segmentation using data from the LiDAR sensor. An efficient hardware architecture is proposed and implemented on an FPGA that can process each LiDAR scan in 17.59 ms, which is much faster than the previous works. Evaluated using Ford and KITTI road detection benchmarks, the proposed solution achieves both high accuracy in performance and real-time processing in speed.

Index Terms—Autonomous vehicle, road segmentation, CNN, LiDAR, FPGA

I. INTRODUCTION

In recent years, we have witnessed a strong increase of research interests on autonomous vehicles [23] [10]. Since the DAPRA Urban Challenge in 2007, automated driving technology has grown rapidly from research experiments to commercial vehicle prototypes owing to the explosive progress in the fields of artificial intelligence and machine learning. As an important task of an automated driving system, it is critical to conduct research on traffic scene perception and its implementations on hardware platforms.

For traffic scene perception, detecting and tracking algorithms are aimed to perceive the surroundings and to set the constraints for planning and control tasks. Based on the object types, the task of traffic scene perception can be classified into three sub-tasks: (1) road perception includes drivable region segmentation and lane detection, (2) object detection/tracking, and (3) traffic sign/signal detection. In road perception, drivable region segmentation scans the front area and searches for the drivable region, while lane detection narrows the region of planning to the ego-lane if lane markers are visible. Object detection and tracking identify the moving objects such as vehicles, pedestrians, cyclists and animals, and measure their locations, dimensions and speed to avoid a collision. Traffic

sign/signal detection looks for traffic signs and traffic lights to perceive additional constraints for planning tasks [8]. As a critical component of an automated driving system, drivable region segmentation provides fundamental knowledge of driving environment. Drivable region segmentation solutions are required to perceive a wide range of view, generate accurate results, and respond in real-time. However, road scenes are complicated. As described in [19], road scenes have three types of diversities: (1) appearance diversity due to changing shapes of lane markers and camera lens distortion, (2) clarity diversity due to occlusions and illumination, and (3) visibility condition diversity due to weather conditions.

Many sensing modalities have been used for drivable region segmentation. Vision modalities [26] [2] [13] are frequently applied on drivable region segmentation for two major reasons: (1) Vision modality is similar to human visual system and most road markers have features in the visual domain, and (2) As a passive sensor, visual camera provides high-resolution data with rich features. By implementing multiple cameras, stereo vision [5] can provide depth information for drivable region segmentation. However, due to the diversity in road scene, it is difficult to design a feature descriptor that handles all visual cases and light conditions. In addition, Shen et al. proposed a series of algorithms to cluster super-pixels that could improve vision based semantic segmentation [30] [29].

Light Detection And Ranging (LiDAR) is another major modality often used by autonomous vehicles. By actively emitting laser beams and measuring the 3D geometry around the vehicle using Time of Flight (ToF), LiDAR can provide a few million geometric points per frame with centimeter accuracy. In addition, LiDAR is not subjected to environmental illumination. However, compared to vision modalities, LiDAR points are sparse and do not contain any visual features employed in traditional vision based algorithms. Several recent works studied traffic scene perception involving LiDAR modality and proposed various schemes for data arrangement, feature extraction and sensor fusion with monocular vision. Much in depth studies are needed on LiDAR data arrangement and feature extraction for accurate and efficient LiDAR based drivable region segmentation.

In the past decades, drivable region segmentation has been studied with different sensors and methodologies. A general solution consists of four components: pre-processing, feature extraction, detection and post-processing. Pre-processing includes noise removal, data sampling and transformation. Feature extraction encodes local features such as color, edge and texture from pre-processed data. Detection applies manually defined or machine learning based models to detect road

This work is partially supported by U.S. NSF Grant CNS-1626236 and by The MathWorks fellowship. The authors are with the Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, Massachusetts 01609, USA. The corresponding author is Xinming Huang (e-mail: xhuang@wpi.edu).

area or lane boundaries. Lastly, post-processing suppresses candidates to provide final results.

In traditional computer vision algorithms, those four steps are totally separated and the extracted features are often describable. However, manually defined features and detectors only work well in normal conditions but cannot handle much variations on the road. Machine learning especially convolutional neural network (CNN) based algorithms combine feature extraction and detection together. Pre-processed data are fed into a well-structured CNN with millions of parameters. Despite that features and detectors are hardly describable visually, machine learning based road perception algorithms have significant advantages in accuracy when compared with traditional computer vision based approach.

For autonomous vehicles, both real-time processing speed and low power consumption are desirable. Graphics processing unit (GPU) devices are popular for parallel processing, but usually consume too much power. Currently only one or two GPU devices can be installed in a vehicle due to the limited power supply. But tens of perception and planning tasks need to be processed on the GPUs simultaneously. Field-programmable gate arrays (FPGAs) are low-power devices that are more suitable for embedded systems. Moreover, an FPGA can be developed as a customized integrated circuit that is able to perform massive parallel processing and data communications on-chip. Hereby, FPGA is our chosen platform that meets both computational capability requirement and power consumption constraint.

In this paper, we present ChipNet as a CNN-based algorithm and its FPGA implementation for real-time LiDAR data processing. The contributions of our work can be summarized as follows: (1) We introduce a new data organizing and sampling method in spherical coordinate that improves the usage of LiDAR points and creates a dense input tensor for CNN. (2) We propose an efficient convolution block for CNN that is both hardware friendly and extendable. (3) The proposed approach of drivable region segmentation results the state-of-art accuracy when evaluated using Ford dataset and KITTI benchmark. We also labelled the Ford dataset for training and evaluation. (4) An efficient and flexible 3D convolution module is designed and implemented on an FPGA, which can achieve real-time processing speed with limited hardware resource and power usage.

The rest of the paper is organized as follows. Section II introduces the related works on road perception task. The proposed drivable region segmentation algorithm is described and its performance on benchmarks are presented in Section III. Section IV presents the FPGA architecture and hardware implementation results. Finally, Section V concludes the paper.

II. RELATED WORK

LiDAR data arrangement: There exists various methods of LiDAR data arrangement on traffic scene perception. In Soquet et al. [33], Alvarez et al. [1], Shinzato et al. [31] and Liu et al. [25], LiDAR point cloud was projected to image

view and manually defined features were applied based on evaluation measurements with image patches. Similar image view was employed by Han et al. [16] and Gu et al. [15] followed by feature extraction using histogram. Gonzalez et al. [14] and Xiao et al. [38] created a dense depth map from point cloud and then combined the map with the camera data for their machine learning based road boundary detector. Similarly, the multi-view method [6] transformed point cloud into both image and top views and then combined with camera data for sensor fusion using a CNN. In addition, VoxelNet [41] and 3D-FCN [24] directly processed sparse LiDAR data in world coordinate using convolutional neural network. LoDNN [3] organized the point cloud into a top view and then fed it into a CNN to generate a heat map representing the possibility of drivable region in each $0.1m \times 0.1m$ cell.

Beside road perception, several research works proposed using CNN for LiDAR-based vehicle detection [41] [6]. To overcome the shortage of training samples, data augmentation and coarse labeling methods were proposed to enlarge the dataset. VoxelNet [41] augmented training data by rotating and translating LiDAR points together with ground truth. StixelNet [13] used LiDAR points to generate coarse labeling automatically for pre-training.

CNN for road perception: Convolutional neural networks have become an active approach for the task of road perception. Starting from Fully Convolutional Network (FCN) [26], various network structures have been proposed to provide accurate road detection and segmentation. SegNet [2] introduced an encoder-decoder scheme to separate feature extractor and detector components. It also added additional connections between the encoder and decoder layers that improved the training of the first few layers closer to the input. Oliveira et al. [27] followed the encoder-decoder scheme and perceived near range and far range in separate branches that resulted an increased accuracy of vision based segmentation. RBNet [7] also followed the encoder-decoder scheme but connected all encoder layer outputs to the decoders. Other works introduced the CNN for salient object detection in images [35] and videos [36]. Most recently, CNN has also been introduced to LiDAR based road segmentation. LoDNN [3], VoxelNet [41] and Multi-view [6] proposed different techniques on LiDAR based perception.

Embedded platforms for road perception: Considering the situations of automated driving or advanced driver assistance system (ADAS), the processing time of road perception algorithms must fulfill the real-time requirement, and thus are often implemented on embedded platforms such as FPGA, ASIC or a mobile CPU/GPU processor. Huval et al. [20] deployed a neural network on Jetson TK1 mobile GPU platform. It detected lane markers based on images and achieved 2.5 Hz running speed. Similarly, the neural network proposed in [34] was able to segment multiple objects including vehicles, pedestrian and pavements at 10 Hz with image resolution of 320p on TX1 GPU platform. Two FPGA based lane detection solutions were proposed in [40] [34] and their processing time were at 60 Hz and 550 Hz, respectively.




Figure 1: A typical LiDAR frame. (a) LiDAR data matrix, (b) LiDAR points projected on corresponding image, and (c) LiDAR points presented on top view

III. ALGORITHM DESIGN

In this work, a hardware friendly and extendable convolutional neural network is proposed to segment drivable region using LiDAR data. In this section, we first introduce the LiDAR data preparation method as pre-processing of the CNN. Next, the proposed network architecture ChipNet is described in detail. Furthermore, we introduce a simulated quantization scheme for CNN that transforms floating-point to fixed-point operations, and thus speeds up the processing on hardware considerably. Finally, a post-processing algorithm is developed to generate a decision map denoting the drivable regions from the CNN output. The proposed solution is evaluated on Ford Campus Vision and LiDAR dataset and KITTI road benchmark. The performance results are presented towards the end of this section.

A. LiDAR data preparation

Typically a LiDAR device places a number of laser scanners vertically and rotates them azimuthally to scan the surrounding obstacles. Suppose a LiDAR device that contains N scanners, measures M points per second and rotates at R rpm, then it generates $\frac{R}{60}$ frames per second with $\frac{60M}{R}$ measure points per frame at an azimuthal resolution of $\frac{360NR}{60M}$. The polar resolution is $\frac{\phi}{N}$ where ϕ denotes the vertical field of view. For example, the HDL-64E LiDAR used in KITTI road benchmark [12] has 64 scan channels and emits 1.33 million points per second. By rotating at 600 rpm it updates 10 frames per second with 0.133 million measurement points per frame at 0.17° azimuthal resolution. By focusing on a 26.90° vertical field of view, the polar resolution is 0.42° . In practice, the LiDAR sensor occasionally generates void measure points when the laser beam emits to a low reflective surface.

Typically, a frame of data generated by the LiDAR modality is a table as shown in Figure 1. In each row, the measurement of a corresponding LiDAR point is listed in four columns, including location coordinates x, y, z of the LiDAR view and laser reflection intensity of the target surface r . By projecting all points to camera view and top-view, as presented in Figure 1, LiDAR point cloud is sparse and has large variations of point density throughout the entire space. Therefore, LiDAR data needs to be organized and re-sampled before being fed to the convolutional neural network.

As mentioned in Section II, there is no unified method to LiDAR point cloud data arrangement and sampling view. Table

I summarizes several research works that organized LiDAR data in different forms. In Table I, we can see that most of them divided the 3D space in Cartesian coordinates, but they sampled the point cloud in different views, such as top view [3] [6], front view [41] or 3D view [24]. We also find a large percentage of LiDAR points are encoded in their region of interest (RoI). However, the organized data as the input to neural networks are sparse, which means that the majority of computations in the first few layers of CNN actually deal with zeros. That is very inefficient from the computational perspective.

Therefore, we propose to organize the LiDAR data in spherical view as if a LiDAR naturally scans the surroundings, as shown in Figure 3. A region of interest is selected in azimuth $[-45^\circ, 45^\circ]$ and all 64 lines of scan points are involved in segmentation. On each line, scan points are grouped by every 0.5° into cells. In total, all scan points in the RoI fall into a 180×64 mesh. We use 0.5° because it is 3 times of LiDAR azimuthal resolution so that in theory at least 2 scan points are grouped in each cell. In practical terms, there are some void scans when the reflect surface is out of range or has low-reflectivity. Input tensor is built in the same width and height as the scan point mesh, but contains 14 feature channels. In each cell, the first 7 features come from the point nearest to the scanner, the next 7 features come from the point furthest away from the scanner. These features include Cartesian coordinates x, y, z , spherical coordinates θ, φ, ρ and the laser reflection intensity r .

In Table I, we compare the LiDAR data preparation with several related works. By sampling LiDAR points in spherical view, our work not only has high LiDAR point usage in RoI but also creates a dense input tensor that improves the accuracy performance and makes the computations in the CNN much more efficient.

B. ChipNet: a hardware friendly and extendable CNN architecture

In this section, we introduce the ChipNet architecture and its simulated quantization algorithm when training on a GPU. The innovations of ChipNet are: (1) we designed a dilated block equivalent to a 5×5 convolutional kernel but saves parameters and calculations, (2) we designed an extendable CNN structure using the dilated block, and (3) we proposed a simulated quantization algorithm to obtain the fixed-point parameters




Figure 2: Data Flow of the proposed LiDAR processing approach

Table I: Summary of LiDAR data organization as input to neural networks

Method	Space subdivision	Sampling view	LiDAR point usage in input tensor	Cell usage in input tensor
VoxelNet [41]	even on Cartesian	front view	98%	1%
3D-FCN [24]	even on Cartesian	3D view	32%	1%
Multi-view (top-view) [6]	even on Cartesian	bird eye view	21%	15%
Multi-view (front-view) [6]	uneven on camera view	camera view	73%	8%
LoDNN [3]	even on Cartesian	bird eye view	24%	15%
ChipNet (ours)	even on Spherical	spherical view	66%	87%

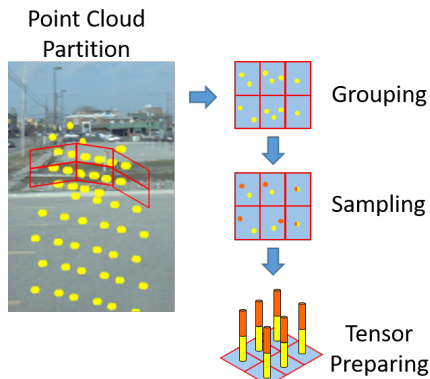


Figure 3: An illustration of the proposed LiDAR data preparation method

for hardware implementations. The network is evaluated using Ford dataset and KITTI road benchmark.

1) *ChipNet convolutional block*: The convolutional block is a key component in ChipNet architecture. Each network block contains three branches. The first one is an identity branch that directly copies the input to the output. As analyzed in [18], identity branch contributes the majority of gradient in back-propagation and decreases the chance of gradient vanishing and explosion during training. The second branch is a 3×3 convolutional layer with 64 channel outputs. The second branch is aimed to encode local features. The third branch is a dilated 3×3 convolutional layer [39] to process features in further pixels but takes less parameters and calculations. As shown in Figure 4, after adding all three branches element-wise, the block equivalent to a 5×5 convolutional layer but has a stable gradient in back-propagation and fewer parameters. Assuming the convolutional layer input size is $180 \times 64 \times 64$ and output size is also $180 \times 64 \times 64$, a ChipNet block contains only 73,856 parameters and requires

802 million multiplications. In comparison, a conventional 5×5 convolutional layer employs 102,464 parameters and requires 1,180 million multiplications. As a results, the proposed ChipNet block reduce parameters by 28% and multiplications by 32%.

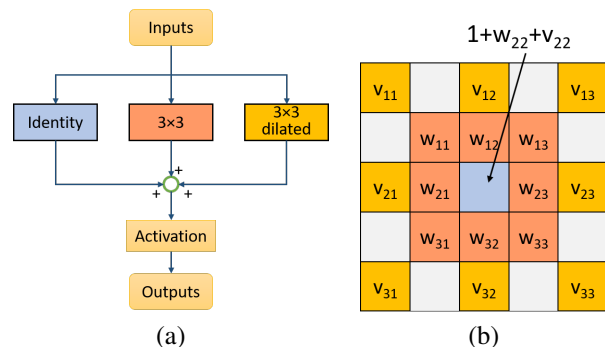


Figure 4: ChipNet convolution block. (a) block architecture and (b) its equivalent 5×5 convolution kernel. The cells in blue, red and brown denote the contributions of corresponding convolutional operations in a conventional 5×5 convolution kernel.

2) *ChipNet network architecture*: The overall CNN architecture of ChipNet is shown in Table II. The first layer is a local feature encoder aimed to encode the input LiDAR data into a 64-channel feature tensor. After encoding, the proposed ChipNet convolution block is instantiated repetitively in the network to perform additional encoding and decoding. Since the input and output of all ChipNet blocks are exactly in the same sizes, the neural network can be conveniently extended deeper by adding more layers. In our work, the ChipNet block is instantiated 10 times as a trade-off between segmentation accuracy and processing latency. For the output layer, a channel-wise mapping is used to generate the final decision map showing the probability of corresponding drivable

Table II: Layer configuration of the ChipNet architecture

Layer	Kernel ($w \times h \times m \times n$)	Input ($w \times h \times m$)	Output ($w \times h \times n$)
Input	—	$180 \times 64 \times 14$	—
Conv Encoder	$5 \times 5 \times 14 \times 64$	$180 \times 64 \times 14$	$180 \times 64 \times 64$
ChipNet Block $\times 10$	$1 \times 1 \times 64 \times 64$ $3 \times 3 \times 64 \times 64$ $3 \times 3 \times 64 \times 64$	$180 \times 64 \times 64$	$180 \times 64 \times 64$
Output	$1 \times 1 \times 64 \times 1$	$180 \times 64 \times 64$	$180 \times 64 \times 1$

regions. Compared to FCN and SegNet, the proposed network is much simpler and more importantly it is extendable. The repetitive network structure is best fitted for hardware reuse in the FPGA design.

3) *Simulated quantization*: Simulated quantization is essential to the training of networks on hardware. Fixed-point variables, weights and operations are widely used in FPGA design, which often utilizes less hardware resources and memories and results higher clock speed, if compared with the floating-point implementations. However, CPU and GPU platforms generally employ floating-point operations that have no quantization error and can generate continuous gradients in the training session. Practically, we can implement CNNs on an FPGA for low-power embedded application. But we still heavily rely on the high-performance GPUs to train the neural networks in order to generate the parameters and weights, since a GPU machine is capable of storing terabytes of training samples and processing hundreds of threads simultaneously.

However, we cannot simply quantize all variables and parameters of a pre-trained neural network from floating-point into fixed-point. Since quantization is a nonlinear operation, the results would not be optimal for the fixed-point neural work. In addition, direct quantization of variables and parameters may result in a loss of gradients. Hereby, we propose a simulated quantization method to train a neural network that can produce the optimal parameters in fixed-point form. This is an essential step to prepare a CNN before implement it on an FPGA.

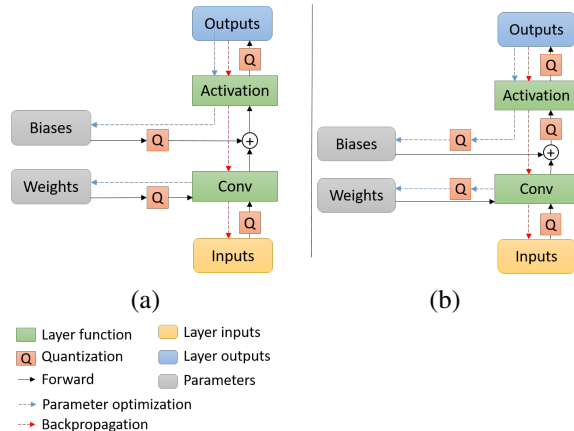


Figure 5: (a) Simulated quantization method as in [21] and (b) the proposed quantization method in this work

Simulated quantization of weights: Quantization that we refer here is not an simple operation of quantizing all weights from floating-point to fixed-point numbers. Additional training is needed to avoid negative impact on the accuracy. At the training stage, however, floating-point weights are preferred because we want to avoid gradient exploding and vanishing. In [21], a simulated quantization approach was proposed in which weights and gradients are stored as floating-point numbers during back-propagation training but the quantized fixed-point numbers are used during forward convolutional operations. The advantage was that the weights and gradients are updated in continuous space so that local optimum due to quantization can be avoided. The disadvantage was that several key functions need to be modified to support this method. However, it is usually difficult to modify, maintain and distribute customized components in a general machine-learning platform such as TensorFlow.

In our work, a new weight regulator is defined and added to the existing network. The regulator is described in *Algorithm 1*. The key innovation is that the regulator quantizes the weights during training and the fixed-point numbers are used during forward operations. Meanwhile, the floating-point weights are also stored in the memory that are used when computing the gradients during back-propagation. However, quantization function is not differentiable. Therefore, we introduce the *StopGradient* function. The *StopGradient* function is a built-in function in TensorFlow that force its gradient to be zero for given input. By applying this function, the gradients are kept the same as of floating point backpropagation, while the weights are quantized. Hence, the proposed quantization algorithm is imported to the TensorFlow platform as a plugin regulator. The proposed quantization algorithm and data flow are shown in Figure 5 in comparison to the simulated quantization in [21].

Algorithm 1 weight quantization

- Data** Weights W , Gradients G
Parameter total_bits= N , fraction_bits= F
- 1: Fraction_scalar $\leftarrow 2^F$
 - 2: Upper_bound $\leftarrow 2^{N-1} - 1$
 - 3: Lower_bound $\leftarrow -2^{N-1}$
 - 4: $\hat{W} \leftarrow W \times \text{Fraction_scalar}$
 - 5: $\hat{W} \leftarrow \text{round}(\hat{W})$
 - 6: $\hat{W} \leftarrow \max(\hat{W}, \text{Lower_bound})$
 - 7: $\hat{W} \leftarrow \min(\hat{W}, \text{Upper_bound})$
 - 8: $\hat{W} \leftarrow \hat{W} / \text{Fraction_scalar}$
 - 9: $W \leftarrow W + \text{StopGradient}(\hat{W} - W)$
 - 10: **return** $\{W, G\}$
-

Simulated quantization of variables: Quantization of variables is similar to the weights. Quantization of variables is implemented as a new activation function so that it can be imported as a custom defined function rather than modifying the body of an existing platform. As described in *Algorithm 1*, if we convert a floating-point number to a N -bit fixed-point number with an F -bit fraction, the operation is to shift to the

left by F bits and then round it to a N -bit integer, followed by shifting back F bits to the right. To minimize negative impact on the back-propagation training, the gradients are all computed in floating-point.

Evaluation of quantization: To evaluate the influence of quantization on accuracy, we first trained the ChipNet in floating point using the Ford training set, and then quantized and fine-tuned using the same training set. Both versions of the ChipNet with and without quantization are evaluated on the same dataset. The result listed in Table III shows that ChipNet quantized to 18 or more bits has similar performance compared to the floating-point model, which indicates that our proposed quantization scheme does not cause accuracy degradation for convolutional neural networks.

C. View of drivable region

The output of network denotes the possibility of drivable region for each cell in spherical view. In post-processing, the output is projected to the top view of a 20-meter wide and 40-meter long area in front of the vehicle. We choose topview in post-processing because it matches the output data format in KITTI benchmark [12], so we can compare our results with others reported in the dataset.

The post-processing algorithm is described in Algorithm 2. Suppose the possibility threshold of a drivable region is set to THR , then the reference point in each column j in the network output $P_{i,j}$ is determined by the nearest LiDAR point in group $\{P|_{col=j, p < THR}\}$. After generating the reference points, a contour of the drivable region becomes a polygon that contains all reference points as vertices.

The post-processing scheme is implemented on CPU using GridMap [11] that is an universal grid map management library. The GridMap library stores map data as Eigen matrix and supports iterators for rectangular, circular, polygonal regions and lines allowing convenient and efficient cell data access. In post-processing, we initialize a grid map instance with a range setting of [6, 46] meters in x -coordinate and [-10, 10] meters in y -coordinate. The resolution is set to 0.05 meter per cell so that the grid map has 800 cells in x -coordinate and 400 cells in y -coordinate. When the post-processing node receives a network output frame, it stores the frame as an Eigen matrix. In Algorithm 2, Step 1-3 is processed on the matrix. In Step 4-5, the contour vertices are imported to a polygon iterator instance and then the drivable region is labeled cell by cell as the polygon iterates. The execution time of post-processing is 5ms per frame on a typical CPU.

The post-processing affects the segmentation accuracy in two parts: (1) After projected to the top view, LiDAR scans are so sparse in far range that there exists distortion between the projected LiDAR boundary and the real drivable region boundary, and (2) Algorithm 2 assumes that the space inside the polygon is drivable while the space outside is not, so there exists an error if the drivable region has holes. We use the KITTI training set to evaluate the effect quantitatively. By projecting the ground truth in Spherical view to top view using the post-processing algorithm and comparing with the top view ground truth, we found that the F1 - score is limited to 95.5% by the post-processing algorithm.

For visualization purpose, we also generate a drivable region map on camera view by applying a similar post-processing procedure.

Algorithm 2 Post-processing of CNN output as segmentation results

Data Input tensor $\mathbf{I}=\{I_{i,j}\}$, **output tensor** $\mathbf{P}=\{P_{i,j}\}$
Parameter threshold THR .
1: $\mathbf{I} \leftarrow Threshold(\mathbf{I} \geq THR)$
2: $\hat{\mathbf{I}} \leftarrow GetLargestConnectedComponent(\mathbf{I})$
3: $\hat{\mathbf{I}} \leftarrow Dilation(\hat{\mathbf{I}}, disk(1))$
4: $\mathbf{B} \leftarrow GetContour(\hat{\mathbf{I}})$
5: $\tilde{\mathbf{B}} \leftarrow ProjectToTargetView(\mathbf{B})$
5: $\tilde{\mathbf{A}} \leftarrow Polygon(\tilde{\mathbf{B}})$
6: **return** $\tilde{\mathbf{A}}$

D. Network training and evaluation

The training platform of ChipNet is a workstation with Xeon 2.4 GHz CPU and NVidia K20 GPU. The software environment is a Python based framework named Keras [9] with TensorFlow 1.4 back end. The input of the network is an $180 \times 64 \times 14$ tensor and the output of the network is an $180 \times 64 \times 1$ tensor. The training speed on the platform is 256 ms per frame. To evaluate the performance of the proposed solution, a subset of the Ford Campus Vision and LiDAR Dataset [28] and the KITTI road benchmark [12] is used for training and testing purposes.

The original Ford Dataset described in [28] contains 3871 frames of LiDAR data recorded with synchronized camera data. The LiDAR data are sampled at 10 Hz. The dataset itself has no labels or annotations, so we created a subset and labeled the drivable region manually. To reduce the overlaps among the consecutive frames, we selected only 1 out of every 5 consecutive frames. Effectively the dataset is downsampled to 2 frames per second. We also removed some off-road samples, such as vehicles on the parking lot, from the dataset since we concentrated on road scenarios. Therefore, we generated a 600-frame subset from the Ford dataset for training and evaluation.

In the subset, the original image is cropped from the size of 1243×1616 to 800×200 resolution that overlaps with the LiDAR point cloud. The data is arranged as described in Section III-A. In order to obtain the LiDAR ground truth, a ray tracing approach is applied as described in Algorithm 3. The projection method from LiDAR coordinate to camera coordinate is described in Algorithm 4. In our labeled subset, each sample includes a $180 \times 64 \times 14$ LiDAR frame, a 800×200 color image, a $180 \times 64 \times 1$ LiDAR ground truth frame and a 800×200 ground truth image. We randomly selected 400 samples for training/validation and the remaining 200 samples for evaluation. Furthermore, we augmented the training samples through rotating the field of view by $(-10, -5, 0, 5, 10)$ degrees from the LiDAR ground truth. Thus, we generated a training set with 2000 samples.

Cross entropy was selected as the loss function and Adam [22] method with default settings was selected as the optimizer. We first trained the network without the quantization plug-in for 30 epochs, at which time the training process converged well. We then fine-tuned the network with the quantization plug-in for 10 epochs to obtain the fixed-point weights. The initial training took 4.5 hours and the fine-tuning took 1.5 hours. For each defined fixed-point bit-length format, we applied the same simulated quantization procedure during fine-tuning. The bit length resulted the least loss is chosen for the FPGA implementation.

Algorithm 3 Ground truth labeling for LiDAR samples

Data Input tensor $L=\{L_{i,j,k}\}$, **Ground truth image** $B=\{B_{i,j}\}$

- 1: $\vec{x}_1 \leftarrow \begin{bmatrix} L_{i,j,1} \\ L_{i,j,2} \\ L_{i,j,3} \end{bmatrix}$, $\vec{x}_2 \leftarrow \begin{bmatrix} L_{i,j,8} \\ L_{i,j,9} \\ L_{i,j,10} \end{bmatrix}$
- 2: $\hat{\vec{x}}_1 \leftarrow Proj(\vec{x}_1)$, $\hat{\vec{x}}_2 \leftarrow Proj(\vec{x}_2)$
- 3: $G_{i,j,1} \leftarrow [B(\hat{\vec{x}}_1) > 0] \times [B(\hat{\vec{x}}_2) > 0]$
- 4: **return** G

Algorithm 4 Projection from LiDAR coordinate to camera coordinate

Data LiDAR point P_{x_i, y_i, z_i, r_i}
Parameter transform matrix $K \in R^{3 \times 4}$

- 1: $\begin{bmatrix} \hat{x}_i \\ \hat{y}_i \\ \hat{z}_i \end{bmatrix} \leftarrow K \begin{bmatrix} x_i \\ y_i \\ z_i \\ 1 \end{bmatrix}$
- 2: $\begin{bmatrix} \hat{x}_i \\ \hat{y}_i \end{bmatrix} \leftarrow \begin{bmatrix} \frac{\hat{x}_i}{\hat{z}_i} \\ \frac{\hat{y}_i}{\hat{z}_i} \end{bmatrix}$
- 3: **return** $\begin{bmatrix} \hat{x}_i \\ \hat{y}_i \end{bmatrix}$

In the testing session, we selected F1 score (F1), average precision (AP), precision (PRE), recall (REC), false positive rate (FPR) and false negative rate (FNR) in image view as the evaluating metrics. The metrics are computed as in (1-4). Table III presents the evaluation results using different bit length of fixed-point quantization. The result shows that the proposed network quantized to 16 or more bits has comparable accuracy to floating-point results, but accuracy drops sharply if quantization is below 16 bits. In our work, 18 bits are selected since it is the best choice supported by the target FPGA platform.

$$Precision = \frac{TP}{TP + FP} \quad (1)$$

$$Recall = \frac{TP}{TP + FN} \quad (2)$$

$$F1 \text{ score} = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall} \quad (3)$$

$$AP = \frac{TP + TN}{TP + FP + TN + FN} \quad (4)$$

We also evaluate our network in KITTI road benchmark [12]. The KITTI vision benchmark suite is a widely used dataset that contains LiDAR, camera, GPS and IMU data. In addition, the vertex transformation from LiDAR coordinate to camera coordinate is provided. The road benchmark in the suite includes 289 training samples and 290 test samples. The point cloud was acquired by a 64-line Velodyne laser scanner and the camera frames were recorded from a Point Grey 1.4 megapixels camera. For better sensor fusion, the LiDAR point cloud is rectified at each time step and the camera frame is cropped to a 375×1242 image. In addition, the data frames from different sensors are synchronized to 10 Hz.

Different from the Ford dataset that evaluates on camera view, the KITTI road benchmark evaluates the segmentation results on top view, in which the result is mapped to a 400×800 image. The mapped image represents the accessibility of the region of 40 meters in the front (from 6 meters to 46 meters) and 10 meters on each side (left and right).

In the training session, we first augmented the dataset through rotating the field of view by $(-10, -8, -6, -4, -2, 0, 2, 4, 6, 8, 10)$ degrees from the corresponding LiDAR ground truth. So, we obtained 3179 samples, among that 3000 randomly selected samples are used for training and the other 179 samples are used for validation. We fine-tuned the network with quantization plug-in for 10 epochs from the weights trained in Ford dataset, and submitted the results to the benchmark online evaluator. The training time was 2.05 hours.

A comparison with several existing results is presented in Table IV. Typical results are shown in Figure 7. The red area denotes the false drivable region (false positive), the blue area denotes the missing driving region (false negative), the green area denotes the correct drivable region (true positive), and the rest area denotes the correct forbidden region (true negative) or don't care region. We also evaluated ChipNet on the front view as in [6] without quantization, which requires significant more run time but results lower accuracy than the spherical view. It implies that LiDAR data arrangement spherical view reveals more features as input to the CNN.

Our proposed approach can provide highly reliable drivable region segmentation with minor distortions around the road boundary. For vehicles on the road, the segmentation boundary matches the ground truth boundary or slightly distorts towards the road center that is safe for automated driving. For the road with sidewalk, the segment boundary matches the ground truth if the sidewalk is above the road surface. However, if the sidewalk is equal or below the road surface, the detected drivable region sometimes extends 1 to 2 meters into the sidewalk, which needs to be improved in future research. In addition, our solution returns accurate drivable regions in poor




Figure 6: Examples of the segmentation results from Ford dataset




Figure 7: Examples of segmentation results from KITTI road dataset

illumination scenarios such as inside tunnels or facing the sun glare. In contrast, vision based solutions rarely work well in those scenarios.

Table III: Performance impact of quantization evaluated on Ford dataset

Name	F1 %	AP %	PRE %	REC %	FPR %	FNR %
ChipNet without Quantization	86.6	94.0	85.5	87.6	14.5	3.5
ChipNet Quantized to 32 bit	86.7	94.2	86.7	86.7	13.3	3.7
ChipNet Quantized to 24 bit	86.9	94.1	85.1	88.8	14.9	3.2
ChipNet Quantized to 18 bit	86.3	94.0	86.0	86.7	14.0	3.7
ChipNet Quantized to 16 bit	86.9	93.6	81.8	90.8	18.2	2.6
ChipNet Quantized to 12 bit	83.7	92.3	78.0	90.3	22.0	2.8
SegNet [2]	89.2	93.6	89.4	88.9	10.5	4.6

IV. HARDWARE ARCHITECTURE

As described in Section III, the LiDAR data after pre-processing has 14 channels and the input data size is 180×64 . After the first layer of convolutional encoding, it becomes a feature map with 64 channels. In the next 10 convolution

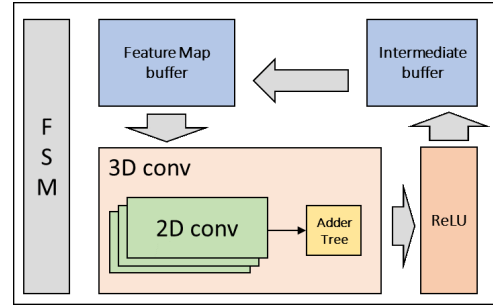


Figure 8: Hardware architecture of ChipNet convolutional neural network

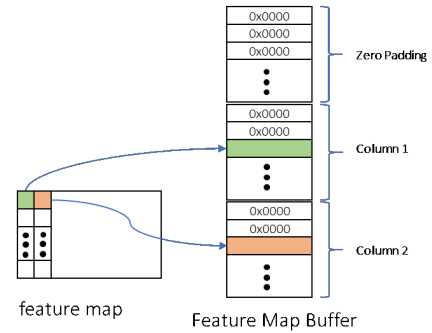


Figure 9: Automatic implementation of zero-padding in hardware

layers, the input and output feature map sizes remain the same as $180 \times 64 \times 64$. The final layer performs the channel-wise mapping that produces an output map of 180×64 , each indicating the possibility of drivable regions. The block diagram of hardware architecture is illustrated in Figure 8. The system consists of a 3D convolution unit, a ReLU block, a feature map buffer and an intermediate buffer. 2D convolution and adder trees are embedded in the 3D convolution block. Since the feature maps in each stage of ChipNet have the same size, this 3D convolution unit is used repetitively. A finite state machine (FSM) is designed to control the iterative processing steps.

A. Zero padding

In order to properly process the information along the boundaries, zero padding must be applied to the feature map produced by the convolution layer output. In our system, a dual-port RAM is implemented for automatic zero-padding. In Figure 9, all memory locations are pre-loaded with zeroes. Pixels of a feature map are written to the corresponding address locations in the feature map buffer. When reading the feature map from the feature map buffer in continuous addresses, data are automatically zero-padded. The RAM functions as the feature map buffer.

Table IV: Comparison with existing results on KITTI road benchmark

Name	F1 %	AP %	PRE %	REC %	FPR %	FNR %	Runtime ms
ChipNet on FPGA (this work)	94.05	88.29	93.57	94.53	3.58	5.47	17.59
LoDNN [3]	94.07	92.03	92.81	95.37	4.07	4.63	18
HybridCRF [38]	90.81	86.01	91.05	90.57	4.90	9.43	1500
LidarHisto [4]	90.67	84.79	93.06	88.41	3.63	11.59	100
MixedCRF [17]	90.59	84.24	89.11	92.13	6.20	7.87	6000
FusedCRF [37]	88.25	79.24	83.62	93.44	10.08	6.56	2000
RES3D-Velo [32]	86.58	78.34	82.63	90.92	10.53	9.08	60
ChipNet on camera view	82.50	86.13	77.37	88.36	14.23	11.64	945

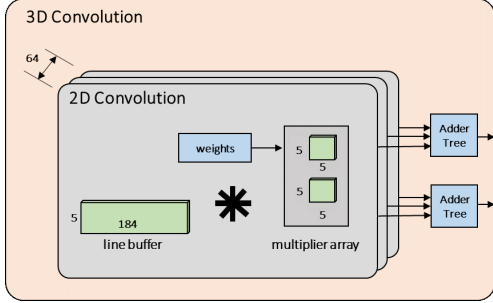


Figure 10: Block diagram of 3D and 2D convolution unit

B. Convolution

As exhibited in Figure 10, the 3D convolution unit contains 64 pieces of 2D convolution slices. Each 2D convolution slice is built with a line buffer and two 5×5 multiplier arrays. The line buffer is designed using shift registers as shown in Figure 11. It outputs a 5×5 window (outlined in red) as the input to the multiplier arrays. The registers in green multiply with the dilated 3×3 convolution kernel, and registers in yellow multiply with the regular 3×3 convolution kernel, and the register at the center multiplies with the coefficient sum as shown in Figure 4.

In each 2D convolution block, the input data are fed from the line buffer to two multiplier arrays, each followed by an adder tree. The 2D convolution block is a pipeline architecture that can process two convolution kernel operations in parallel. Since each 2D convolution operation has 64 convolution kernels, the same feature map is reloaded and processed for 32 times. All weights are stored in on-chip memory to avoid the latency of off-chip memory access. The ReLU block is implemented by a comparator and a multiplexer. If the input value is larger than 0, it outputs the original value. Otherwise the ReLU block outputs 0.

C. FSM Controller

Since the multiplier array in Figure 10 consumes a large number of DSP slices on FPGA, reusing it for each convolution layer is a key consideration in the control logic design. Thus, a cascaded finite state machine (FSM) is deployed to control the iterative process. As shown in Figure 10, the 3D convolution unit can perform 2 kernel operations in parallel. From Table II, each layer requires 64 convolutional kernel

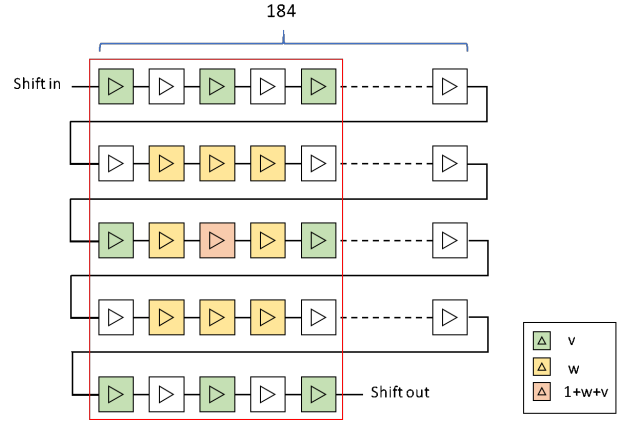


Figure 11: Block diagram of the line buffer unit

operations. So, the inner FSM controls the 3D convolution unit to perform the same operations 32 times, each with different input of feature maps, while the outer FSM controls the order of layers.

As is shown in Figure 8, during each layer of convolution, the outer FSM first loads the input feature map into feature map buffer. Meanwhile, the inner FSM starts to feed feature map into the 3D convolution unit. The intermediate feature maps are stored in intermediate buffer. When the inner FSM completes the convolution of one layer, the outer FSM moves the data from intermediate buffer to feature map buffer, and then it starts the convolution of the next layer.

D. Implementation Results

The target hardware platform is Xilinx UltraScale XCKU115 FPGA. An integrated test system is demonstrated in Figure 12.

The LiDAR frames are transmitted into PC at 10 Hz via UDP protocol. For each LiDAR frame, the PC pre-processes it and sends an 18-bit feature map to the ChipNet neural network in the FPGA. The feature map size is $64 \times 180 \times 14$. The parameters are ported using MATLAB HDL coder. System clock frequency is set to 350 MHz. Each convolution block takes about 12,512 clock cycles to generate 2 feature maps. The total processing time of this CNN architecture is about 12.59 ms. Since normally LiDAR point cloud frame rate is 10 Hz, this FPGA implementation fulfills the requirement of real-time LiDAR data processing. When running ChipNet in




Figure 12: The overall system architecture with a LiDAR and FPGA accelerator

software on the Intel Core i5-5200U CPU, the processing time is 549 ms; when running ChipNet using the NVidia K20 GPU, the processing time is 162 ms. Thus, the FPGA implementation gains $43\times$ speed up over CPU and $13\times$ speed up over GPU.

As mentioned earlier, there are few FPGA implementations of LiDAR processing using CNN at this time, performance and efficiency comparison with similar works on FPGAs is not available.

The resource usage of our proposed neural network is listed in Table V. The total power consumption of this design is 12.594 W, estimated by from Xilinx Vivado 2017.2 power analyzer using post-implementation simulation .saif (Switching Activity Interchange Format) file. As illustrated in Table VI, it consists of dynamic power 9.747W and static power 2.848W. We notice that most of the power is consumed by the on-chip memories since they are always enabled except in idle mode. We emphasize that the proposed FPGA solution takes only 11.8% power consumption of an NVidia K20 GPU, which is 107 W.

Table V: Resource usage on the FPGA implementation of ChipNet

FPGA Resource	Used	Available	Utilization
Slice Registers	33530	1326720	2.53%
Slice LUTs	38082	663360	5.74%
Block RAMs	1543	2160	71.44%
DSPs	3072	5520	55.65%

Table VI: Power estimation of FPGA design

Power Type	Item	Power Consumed
Dynamic	Logic	0.212 W
	BRAM	8.609 W
	DSP	0.006 W
	MMCM	0.26 W
Satistic		2.848 W

V. CONCLUSIONS


In this paper, the problem of drivable region segmentation is framed as a semantic segmentation task by processing real-time LiDAR data using a convolutional neural network on an FPGA. The LiDAR data is organized in spherical view and sampled to a dense input tensor during pre-processing. An efficient and extendable CNN architecture namely ChipNet

is proposed as the main processor. A reusable and efficient 3D convolution block is designed for FPGA implementation. The proposed approach is trained using Ford dataset and the KITTI benchmarks. Evaluations show the proposed LiDAR processing algorithm can achieve state-of-art performance in accuracy and also real-time processing in speed on the FPGA. However, the FPGA implementation still consumes a large amount of on-chip memory. For future work, we will consider recurrent neural network for spatial-sequence decoding that may reduce the on-chip memory usage. We also notice during benchmark evaluation that sidewalk and railway are the main causes of false positives. Sensor fusion of LiDAR and camera data will be considered to further improve the accuracy.


REFERENCES

- [1] José M Álvarez Alvarez and Antonio M Lopez. Road detection based on illuminant invariance. *IEEE Transactions on Intelligent Transportation Systems*, 12(1):184–193, 2011.
- [2] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.
- [3] Luca Caltagirone, Samuel Scheidegger, Lennart Svensson, and Mattias Wahde. Fast lidar-based road detection using fully convolutional neural networks. In *Intelligent Vehicles Symposium (IV), 2017 IEEE*, pages 1019–1024. IEEE, 2017.
- [4] Liang Chen, Jian Yang, and Hui Kong. Lidar-histogram for fast road and obstacle detection. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 1343–1348. IEEE, 2017.
- [5] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In *Advances in Neural Information Processing Systems*, pages 424–432, 2015.
- [6] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *IEEE CVPR*, volume 1, page 3, 2017.
- [7] Zhe Chen and Zijing Chen. Rbnet: A deep neural network for unified road and road boundary detection. In *International Conference on Neural Information Processing*, pages 677–687. Springer, 2017.
- [8] Zhilu Chen and Xinming Huang. Accurate and reliable detection of traffic lights using multiclass learning and multiobject tracking. *IEEE Intelligent Transportation Systems Magazine*, 8(4):28–42, 2016.
- [9] Francois Chollet et al. Keras. <https://keras.io>, 2015.
- [10] A. Cosgun, L. Ma, J. Chiu, J. Huang, M. Demir, A. M. Añon, T. Lian, H. Tafish, and S. Al-Stouhi. Towards full automated drive in urban environments: A demonstration in gomentum station, california. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1811–1818, June 2017.
- [11] Peter Fankhauser and Marco Hutter. A universal grid map library: Implementation and use case for rough terrain navigation. chapter 5. Springer, 2016.
- [12] Jannik Fritsch, Tobias Kuehnl, and Andreas Geiger. A new performance measure and evaluation benchmark for road detection algorithms. In *International Conference on Intelligent Transportation Systems (ITSC)*, 2013.
- [13] Noa Garnett, Shai Silberstein, Shaul Oron, Ethan Fetaya, Uri Verner, Ariel Ayash, Vlad Goldner, Rafi Cohen, Kobi Horn, and Dan Levi. Real-time category-based and general obstacle detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 198–205, 2017.
- [14] Alejandro González, Gabriel Villalonga, Jiaolong Xu, David Vázquez, Jaime Amores, and Antonio M López. Multiview random forest of local experts combining rgb and lidar data for pedestrian detection. In *Intelligent Vehicles Symposium (IV), 2015 IEEE*, pages 356–361. IEEE, 2015.
- [15] S. Gu, Y. Zhang, J. Yang, and H. Kong. Lidar-based urban road detection by histograms of normalized inverse depths and line scanning. In *2017 European Conference on Mobile Robots (ECMR)*, pages 1–6, Sept 2017.
- [16] Xiaofeng Han, Huan Wang, Jianfeng Lu, and Chunxia Zhao. Road detection based on the fusion of lidar and image data. *International Journal of Advanced Robotic Systems*, 14(6):1729881417738102, 2017.

- [17] Xiaofeng Han, Huan Wang, Jianfeng Lu, and Chunxia Zhao. Road detection based on the fusion of lidar and image data. *14(6):172988141773810*, 2017.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Aharon Bar Hillel, Ronen Lerner, Dan Levi, and Guy Raz. Recent progress in road and lane detection: a survey. *Machine vision and applications*, 25(3):727–745, 2014.
- [20] Brody Huval, Tao Wang, Sameep Tandon, Jeff Kiske, Will Song, Joel Pazhayampallil, Mykhaylo Andriluka, Pranav Rajpurkar, Toki Migimatsu, Royce Cheng-Yue, et al. An empirical evaluation of deep learning on highway driving. *arXiv preprint arXiv:1504.01716*, 2015.
- [21] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. *arXiv preprint arXiv:1712.05877*, 2017.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J. Z. Kolter, D. Langer, O. Pink, V. Pratt, M. Sokolsky, G. Stanek, D. Stavens, A. Teichman, M. Werling, and S. Thrun. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 163–168, June 2011.
- [24] Bo Li. 3d fully convolutional network for vehicle detection in point cloud. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*, pages 1513–1518. IEEE, 2017.
- [25] Xiaolong Liu and Zhidong Deng. A graph-based nonparametric drivable road region segmentation approach for driverless car based on lidar data. In *Proceedings of the 2015 Chinese Intelligent Automation Conference*, pages 431–439. Springer, 2015.
- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [27] Gabriel Leivas Oliveira, Wolfram Burgard, and Thomas Brox. Efficient deep methods for monocular road segmentation. In *IEEE/RSJ international conference on intelligent robots and systems (IROS 2016)*, 2016.
- [28] Gaurav Pandey, James R McBride, and Ryan M Eustice. Ford campus vision and lidar data set. *The International Journal of Robotics Research*, 30(13):1543–1552, 2011.
- [29] Jianbing Shen, Yunfan Du, Wenguan Wang, and Xuelong Li. Lazy random walks for superpixel segmentation. *IEEE Transactions on Image Processing*, 23(4):1451–1462, 2014.
- [30] Jianbing Shen, Jianteng Peng, Xingping Dong, Ling Shao, and Fatih Porikli. Higher order energies for image segmentation. *IEEE Transactions on Image Processing*, 26(10):4911–4922, 2017.
- [31] Patrick Y Shinzato, Diego Gomes, and Denis F Wolf. Road estimation with sparse 3d points from stereo data. In *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*, pages 1688–1693. IEEE, 2014.
- [32] Patrick Yuri Shinzato, Denis Fernando Wolf, and Christoph Stiller. Road terrain detection: Avoiding common obstacle detection assumptions using sensor fusion. In *Intelligent Vehicles Symposium (IV)*, 2014.
- [33] Nicolas Soquet, Didier Aubert, and Nicolas Hautiere. Road segmentation supervised by an extended v-disparity algorithm for autonomous navigation. In *Intelligent Vehicles Symposium, 2007 IEEE*, pages 160–165. IEEE, 2007.
- [34] Wei Wang and Xinming Huang. An fpga co-processor for adaptive lane departure warning system. In *Circuits and Systems (ISCAS), 2013 IEEE International Symposium on*, pages 1380–1383. IEEE, 2013.
- [35] Wenguan Wang, Jianbing Shen, and Haibin Ling. A deep network solution for attention and aesthetics aware photo cropping. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [36] Wenguan Wang, Jianbing Shen, and Ling Shao. Video salient object detection via fully convolutional networks. *IEEE Transactions on Image Processing*, 27(1):38–49, 2018.
- [37] Liang Xiao, Bin Dai, Daxue Liu, Tingbo Hu, and Tao Wu. Crf based road detection with multi-sensor fusion. In *Intelligent Vehicles Symposium (IV)*, 2015.
- [38] Liang Xiao, Ruili Wang, Bin Dai, Yuqiang Fang, Daxue Liu, and Tao Wu. Hybrid Conditional Random Field based Camera-LIDAR Fusion for Road Detection. *Information Sciences.*, (XX):1–11, in press.
- [39] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016.
- [40] Jin Zhao, Bingqian Xie, and Xinming Huang. Real-time lane departure and front collision warning system on an fpga. In *High Performance Extreme Computing Conference (HPEC), 2014 IEEE*, pages 1–5. IEEE, 2014.
- [41] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. *CoRR*, abs/1711.06396, 2017.



Yecheng Lyu (S'17) received his B.S. degree from Wuhan University, China in 2012 and M.S. degree from Worcester Polytechnic Institute, USA in 2015 where he is currently a Ph.D student working on autonomous vehicles. His current research interest is sensor fusion, autonomous vehicle perception and deep learning.



Lin Bai (S'18) received the B.S degree in integrated circuits design and integrated system from University of Electronic Science and Technology of China in 2009 and M.S. degree in electrical engineering and information technology from Swiss Federal Institute of Technology Zurich in 2012. Previously he was an FPGA engineer in industry. He is currently working towards the Ph.D. degree at Worcester Polytechnic Institute, USA. His current research interest is hardware acceleration of deep learning algorithms on FPGA and ASIC.



Xinming Huang (M'01–SM'09) received the Ph.D. degree in electrical engineering from Virginia Tech in 2001. Since 2006, he has been a faculty in the Department of Electrical and Computer Engineering at Worcester Polytechnic Institute (WPI), where he is currently a chair professor. Previously he was a Member of Technical Staffs with the Bell Labs of Lucent Technologies. His main research interests are in the areas of circuits and systems, with emphasis on autonomous vehicles, deep learning, IoT and wireless communications.