

K-Lane: Lidar Lane Dataset and Benchmark for Urban Roads and Highways

Dong-Hee Paek^{1*} Seung-Hyung Kong^{1†} Kevin Tirta Wijaya²
¹CCS Graduate School of Mobility ²Robotics Program
Korea Advanced Institute of Science and Technology, Daejeon, Korea
{donghee.paek, skong, kevin.tirta}@kaist.ac.kr

Abstract

Lane detection is a critical function for autonomous driving. With the recent development of deep learning and the publication of camera lane datasets and benchmarks, camera lane detection networks (CLDNs) have been remarkably developed. Unfortunately, CLDNs rely on camera images which are often distorted near the vanishing line and prone to poor lighting condition. This is in contrast with Lidar lane detection networks (LLDNs), which can directly extract the lane lines on the bird’s eye view (BEV) for motion planning and operate robustly under various lighting conditions. However, LLDNs have not been actively studied, mostly due to the absence of large public lidar lane datasets. In this paper, we introduce KAIST-Lane (K-Lane), the world’s first and the largest public urban road and highway lane dataset for Lidar. K-Lane has more than 15K frames and contains annotations of up to six lanes under various road and traffic conditions, e.g., occluded roads of multiple occlusion levels, roads at day and night times, merging (converging and diverging) and curved lanes. We also provide baseline networks we term Lidar lane detection networks utilizing global feature correlator (LLDN-GFC). LLDN-GFC exploits the spatial characteristics of lane lines on the point cloud, which are sparse, thin, and stretched along the entire ground plane of the point cloud. From experimental results, LLDN-GFC achieves the state-of-the-art performance with an F1-score of 82.1%, on the K-Lane. Moreover, LLDN-GFC shows strong performance under various lighting conditions, which is unlike CLDNs, and also robust even in the case of severe occlusions, unlike LLDNs using the conventional CNN. The K-Lane, LLDN-GFC training code, pre-trained models, and complete development kits including evaluation, visualization and annotation tools are available at <https://github.com/kaist-avelab/k-lane>.

1. Introduction

Autonomous driving depends on a number of critical functions that are realized with state-of-the-art (SOTA) technologies. Among them, lane detection function is to detect the accurate location and curvature of the ego lane and neighboring lanes, and provide necessary input to the path planning function. Therefore, the lane detection function should be robust to various conditions (e.g., night, day times) and challenging situations (e.g., lane line occlusion). However, the conventional lane detection techniques based on image processing are vulnerable to situations when lane lines are partially missing or occluded, because the techniques rely on heuristic methods such as denoising, edge detection, and line fitting with detected edges [2, 3, 20].

Recently, lane detection [13, 16, 23] has been largely improved due to the deep learning. When a large dataset with accurate label is available for training, deep learning networks can produce high-quality predictions that are almost indistinguishable to the ground truths. This is the case for camera lane detection networks (CLDNs) [13, 23], which show superior performances compared to the conventional (i.e., heuristic) lane detection techniques when given abundant training samples from public datasets such as CULane [16] and TuSimple [25].

However, CLDNs still have a few inherent problems. First, cameras suffer from poor lighting conditions, such as low or harsh lights [16]. Second, it is often necessary to project front camera images into 2-dimensional (2D) bird’s eye view (BEV) for motion planning, which often causes lane line distortions [1]. For example, BEV-projection of the detected lanes near the vanishing line of a front camera image [27] may result in inaccurate and distorted lane lines so that the motion planning should be limited to a shorter distance.

On the other hand, Lidar has multiple advantages over the camera in the lane detection; lane detection from a Lidar point cloud does not suffer from the distortion in the BEV-projection and is not affected by lighting conditions. However, there have been a little studies introduced in the literature, mostly because there are not enough public dataset

*equal contribution

†corresponding author

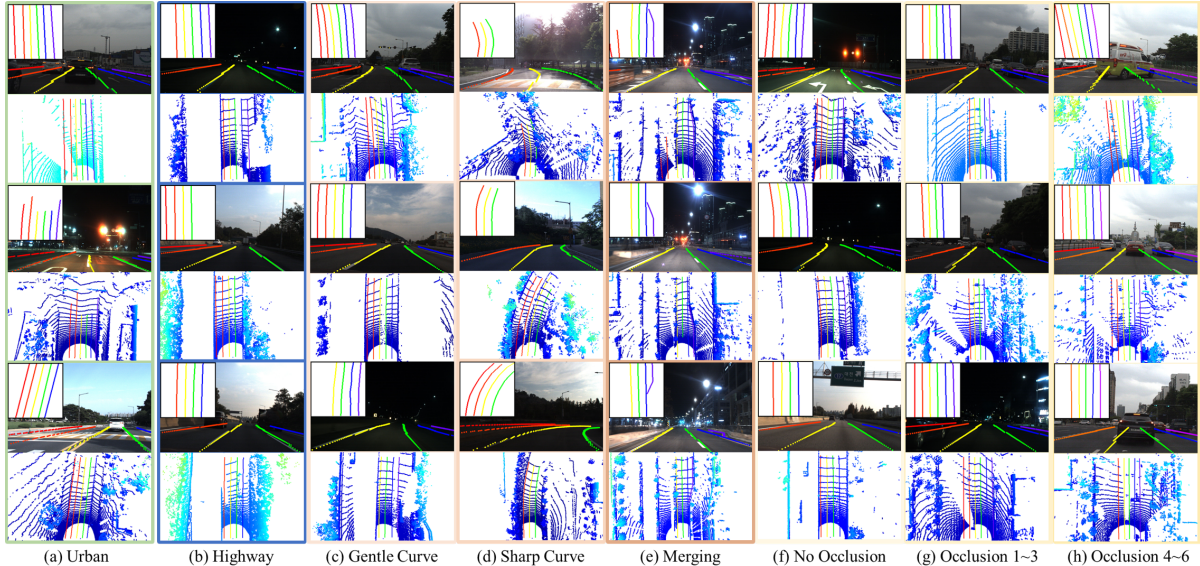


Figure 1. Examples of frames under various conditions for K-Lane, where each column shows one of the conditions: Each column consists of a total of three rows. Each row shows an upper plot for the projection of true BEV labels into the front view image with true BEV labels in the upper left corner and a lower plot for the lane labels on top of the BEV point cloud, respectively.

and benchmark for Lidar lane detection network (LLDN).

In this paper, we introduce KAIST-Lane (K-Lane) dataset, the world’s first and the largest open Lidar lane dataset for Lidar lane detection in urban roads and highways. We also provide an easy-to-use development kits (devkits) for the training, evaluation, dataset development, and visualization. K-Lane has more than 15K annotated frames, and contains a maximum of six lanes under various road and traffic conditions, such as roads at night and day times, merging (converging and diverging) and curved lanes as shown in Fig. 1. Each annotation consists of the lane lines segmentation label, driving condition, lane shape, and occlusion level. As such, the performance of developed LLDNs in different challenging conditions can be evaluated easily, e.g., when driving in the night time or with significant measurement loss due to high occlusions. The segmentation label is accurately annotated with one pixel width on the BEV image, which translates to a $4\text{cm} \times 4\text{cm}$ area in the real world. The label consists of a class id, which represent the relative position of the lane line to the ego-lane. This enables the LLDNs to be trained directly to infer the location of ego-lane, which is crucial for the motion planning. Moreover, as shown in Fig. 1, front camera images have been elaborately calibrated with Lidar point clouds, enabling intuitive visualization and may pave the way for further lane detection studies using multi-modal sensors (e.g., sensor fusion).

To demonstrate the viability of developing LLDNs with K-Lane, we propose a baseline model, Lidar lane detection network utilizing global feature correlator (LLDN-GFC),

which fully exploits the spatial characteristics of the lane lines in point clouds. This is in contrast to most of the CNN-based LLDNs introduced in the literature [1, 14], which are mostly a modification of the CNN-based CLDNs developed for camera images. We observe that the CNN-based LLDNs are not suitable for detecting lane lines in Lidar point cloud. For example, lane lines on the front view image have decreasing thickness with the distance from the ego vehicle and are heading to the same vanishing point (on a straight road), whereas lane lines in a BEV image have a constant thickness and stretch long in parallel over the entire BEV image. These spatial characteristics of the lane lines in Lidar point cloud are not appropriately exploited by the CNN-based lane detection networks, in contrast to our proposed LLDN-GFC. The proposed LLDN-GFC can be implemented with Transformer [4] and Mixer [24] blocks to perform an effective global feature correlation for lane lines. In the experimental results, we show that the proposed baseline achieves a superior performance than LLDNs using the conventional CNN. The contribution of this paper can be summarized as;

- K-Lane: we introduce the world first and the largest (15382 frames) public Lidar lane dataset for urban roads and highways under various conditions and scenarios.
- We also provide a complete devkits for training, evaluation, annotation, and visualization.
- We show that lane lines in the Lidar point cloud have a special characteristics not found in traditional RGB images, and provide appropriate baseline net-

Dataset	Year	Sensor	Num. Lanes	Class Info.	Num. Frames	Public	Road Type
TuSimple [25]	2017	Camera	5	No	6408	Yes	Highway
CULane [16]	2018	Camera	4	Yes	133235	Yes	Urban & Highway
DeepLane [1]	2019	Lidar, Camera	n/a	No	55168	No	Urban & Highway
RoadNet [14]	2020	Lidar	n/a	Yes	5200	No	Highway
K-Lane (ours)	2022	Lidar, Camera	6	Yes	15382	Yes	Urban & Highway

Table 1. Comparison of previous lane datasets and ours (n/a denotes not available).

work we term LLDN-GFC, which significantly outperforms LLDNs using the conventional CNN in the F1-score.

This paper is organized as follows. Section 2 introduces prior studies related to this paper and the topic of this paper, Section 3 introduces K-Lane dataset, and the proposed baseline, LLDN-GFC. Section 4 shows the experiment setup and results. We draw our conclusion in Section 5, and introduce more information for both dataset and baseline, such as detailed network structure of LLDN-GFC in the Appendix.

2. Related Works

Lane Detection Datasets and Benchmarks. Lane detection with data-driven approaches such as deep learning has seen tremendous advancements in recent years. One key enabler towards such advancements is the availability of large public lane dataset, as shown in Table 1. TuSimple [25] is one of the earliest publicly available camera-based lane dataset. It has 6,408 number of frames collected in the highway during the day. The dataset is further divided into 3,626 frames for training, 358 frames for validation, and 2,782 frames for testing. CULane [16] introduces a more diverse and challenging camera-based lane dataset, with 133,235 number of frames divided into 88,880 frames for training, 9,675 frames for validation, and 34,680 frames for testing. CULane provides diverse driving conditions, both in urban and highway environments, in the day and night, and with various road structures. Comparing to the vibrant field of camera-based lane detection, Lidar lane detection dataset has not been explored as much. One of the earliest Lidar lane dataset is DeepLane [1], which contains 55,168 frames of Lidar and camera data collected in both urban and highway environments. Another dataset, RoadNet [14], consists of 5,200 frames of Lidar data collected only in the highway environments. Unfortunately, both datasets are not public, as such not many derivative works on Lidar lane detection have been conducted. In contrast, our proposed dataset, K-Lane, contains 15,382 frames of Lidar and camera data, collected in both urban and highway environments. As we make K-Lane public, we pave the way for a new research direction in lane detection approaches using Lidar.

Lane Detection Networks for Camera. As labeled cam-

era dataset [16] for various roads become available, there have been a significant advancement in the CLDNs. Compared with the early rule-based techniques [2, 3], CLDNs are more adaptive to various road environments. In these techniques, lanes are predicted based on local features extracted by CNN [6], and the performance is improved with lane detection heads that exploit the features of lane lines. For example, Qin et al. [18] proposes a row-wise detection-based network that divides the entire image into grids, and recognize lanes from each row of grids. Liu et al. [13] proposes a two-stage lane detection network that combines the conditional convolution [30] with the row-wise detection in the detection head and achieves SOTA performance in several datasets. However, CLDNs have some inherent problems. In the CULane benchmark, most of the CLDNs show significant performance drop (about 20%) for night time and dazzling light conditions from their daytime performance [13, 18].

Early Lane Detection Techniques for Lidar. In early studies, lane points are detected by thresholding the measured intensity (or reflectivity). Lindner et al. [12] uses a fixed polar grid map to store point intensities and filter the lane candidates with thresholding along azimuth angles. Hernandez et al. [7] introduces a clustering approach, where the filtered lane points are clustered using DBSCAN [5]. However, these heuristic techniques rely on pre-defined thresholding parameters, and, therefore, it is not very adaptive to different environments.

Lane Detection Networks for Lidar. Deep learning-based lane detection studies for Lidar have not been actively conducted due to the absence of large open datasets, and only some studies with their private Lidar datasets are introduced in the literature. Bai et al. [1] proposes an LLDN that combines 2D BEV images developed with the Lidar point cloud and the front camera image for lane detection. And Martinek et al. [14] proposes a CNN-based LLDN that uses BEV images from point clouds to detect ego-lanes, and tests the network in an uncrowded highway.

Self-Attention in Vision. Self-attention is a scheme to lead a neural network to pay more attention to the patches of the input image, between which there is high correlation score. Convolutional block attention module (CBAM) [28] introduces per-channel and per-space self-attention mechanisms by adding MLP (Multilayer Perceptron) and convolu-

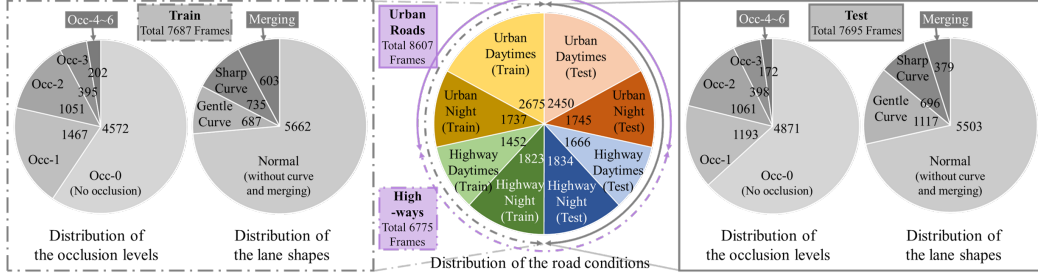


Figure 2. Data distribution (in number of data frames) of the K-Lane. The pie-chart in the middle shows data distribution in road types (i.e., urban roads or highways), time (i.e., daytimes or night), and usage (i.e., training or test). The four pie-charts on the left (for training) and right (for test) shows data distribution with respect to the six levels of lane occlusion (from zero occluded lanes to six occluded lanes) and that with respect to the lane shape (gentle curve, sharp curve and merging).

tional operations, respectively, to the traditional CNN-based feature extractor. Since Transformer [26] shows significant improvement in the self-attention mechanism by applying three independent MLPs for query, key, and value (i.e., Transformer block), it has been used actively for images and point cloud. As an example, ViT (Vision Transformer) [4] greatly improves image classification performance using Transformer, where ViT divides input image into unit patches and applies Transformer encoder to each patch for image classification. However, ViT employs three independent MLPs for each attention mechanism, for which high computational cost and large model size are inevitable. On the other hand, MLP-Mixer [24] implements the attention mechanism with a simple MLP scheme (i.e., Mixer block), which results in a fast inference with a small model size and achieves comparable performance to ViT.

3. K-Lane and LLDN-GFC

In this section, we introduce K-Lane dataset, benchmark, and the proposed baseline, LLDN-GFC.

3.1. K-Lane

K-Lane is the first large open LiDAR lane dataset that consists of Lidar point clouds and their corresponding RGB images for urban roads and highways under various conditions and scenarios as shown in Fig. 1.

Data Distribution. As shown in Fig. 2, there are a total of 15382 data frames, divided into 7687 frames for training and 7695 frames for testing. Each set contains various road conditions and challenging scenarios including (a) different lighting conditions such as day and night times, (b) crowded traffic with lane occlusions by other vehicles, and (c) merging (converging, diverging) and curved lanes, which are further classified into gentle curves and sharp curves. Note that K-Lane has maximum six lanes and occlusions are divided into six levels representing 0, 1, 2, 3, and 4~6 occluded lanes. The benchmark kit provides evaluation tools for calculating the metrics per each condition, and given condi-

tions are annotated for each frame under a clear criterion, which are described in Appendix A.

Sensor Suite. The K-Lane is collected using Ouster OS2-64 Lidar sensor [15] that has 64 channels with a maximum range of 240m, placed on the roof of the vehicle, and a front camera that has 1920×1200 resolution, as shown in Fig. 3-a. Front camera images have been carefully calibrated with Lidar point clouds, which makes it easy to visualize and may enable further lane detection studies with multi-modal sensors.

Dataset Development. The ground truth labels are produced by projecting the Lidar point cloud into BEV, thresholding the intensity measurements to extract keypoints (i.e., candidates of lane lines), and drawing one-pixel-wide line for each lane as shown in Fig 3-b. As such, high resolution and accurate labels are produced, which is critical for deep learning-based methods.

Metrics. To standardize the evaluation of the network being developed, we choose to use the F1-score metric for both confidence and classification, which evaluates per-pixel presence of lane and per-pixel correct classification of the lane, respectively. The F1 metric represents a harmonic mean between precision and recall, and can be expressed as

$$F1 = \frac{1}{\frac{1}{\text{Precision}} + \frac{1}{\text{Recall}}} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2 \cdot \text{TP}}{2 \cdot \text{TP} + \text{FP} + \text{FN}}, \quad (1)$$

where TP, FP, and FN are the numbers of true positives, false positives, and false negatives, of the output of the detection head, respectively. Since the width of a lane line in the label is only one pixel-wide, we allow up to one pixel deviation between the prediction and the label. This is comparable to the evaluation metric used in the CULane [16] dataset, where a lane line label is 30-pixels wide and a true positive is counted when the prediction and the ground truth have at least an IoU of 0.5.

To formally describe the evaluation metric, let $\mathbf{x}^{conf} \in \mathbb{R}^{M \times N}$ be the confidence map label with M number of rows and N number of columns, with $x_{m,n}^{conf} \in \{0, 1\}$. Furthermore, let $\hat{\mathbf{x}}^{conf} \in \mathbb{R}^{M \times N}$ be the confidence map predic-

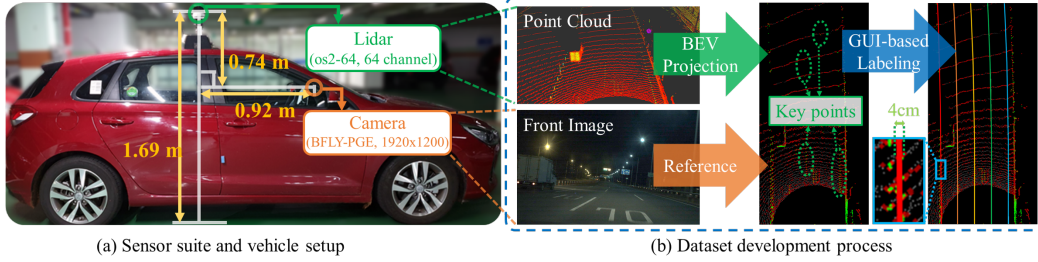


Figure 3. Sensor suite, vehicle setup, and dataset development process of K-Lane.

tion with M number of rows and N number of columns, with $x_{m,n}^{conf} \in [0, 1]$. Additionally, we define a grid neighborhood centered at the grid $x_{m,n}$ as a set of grids $\{x_{i,j} | i = \{m-1, m, m+1\}, j = \{n-1, n, n+1\}\}$.

Suppose that a thresholding operation is applied to the confidence map prediction such that

$$\hat{x}_{m,n}^{conf,thr} = \begin{cases} 1 & \hat{x}_{m,n}^{conf} > \sigma^{conf} \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where the σ_{conf} is the confidence threshold for a prediction to be considered as a lane point. In our evaluation metric, a true positive occurs if for a positive prediction (pixel value equals to 1) at $\hat{x}_{m,n}^{conf,thr}$, there exists at least one positive label at the grid neighborhood centered around $x_{m,n}^{conf}$. Conversely, if there is no positive label at the grid neighborhood, the prediction counts as a false positive. A false negative occurs if for a positive label at $x_{m,n}^{conf}$, there is no positive prediction at the grid neighborhood centered around $\hat{x}_{m,n}^{conf,thr}$.

For classification predictions, we transform the classification map label into a one-hot-encoding label. In addition, we also transform the classification prediction map into a one-hot-encoding prediction where the class with the highest probability is assigned a value of 1 and the rest are set as 0. The true positives, false positives, and false negatives of the classification predictions can be calculated with the previously mentioned procedure, accumulated for every possible classes.

The F1-score on classification is an evaluation of networks based on both lane line localization and lane class prediction. Therefore, F1-score on classification is a strict evaluation metric and, as a result, performance degradation can be found for all models compared to the confidence prediction, as shown in Table 2.

Complete Devkits. Additionally, we provide a comprehensive devkits of K-Lane that includes training, evaluation, dataset development, and visualization. In particular, data development tools, such as labeling and annotation tools, are provided through the Graphic User Interface (GUI) for easy-to-use. This enables the research community to readily increase the dataset regardless of the Lidar sensor models, and thus to activate areas of LLDN with diverse datasets

and benchmarks, as well as CLDN. Appendix A presents a full description of all of the specifics.

Summary. In summary, compared to the conventional lane detection datasets, K-Lane has multiple advantages; (1) K-Lane is collected in urban roads and highways under various conditions and scenarios as stated above, while TuSimple [25] and RoadNet [14] include only highway, (2) K-Lane distinguishes lane classes and labels with precise lane location (pixel level), whereas TuSimple [25] and DeepLane [1] have labels without distinction between lane classes (3) K-Lane has larger number of labeled lanes (e.g., maximum 6 lanes), while TuSimple [25], and CULane [16] have only up to 5 and 4 lanes, respectively, and (4) Above all, among Lidar lane datasets, K-Lane is the only publicly available dataset, which allows more studies on Lidar-based lane detection to be conducted. In addition, the well-calibrated camera images may also be used in future works for multi-modal lane detection.

3.2. LLDN-GFC

In this section, we focus on the overall structure and necessity of our baseline for LLDNs while the details, such as the exact neural network structure, functions (i.e., (1)~(5) in Fig. 4), and mathematical expression for loss, are described in Appendix B. As shown in Fig. 4, the proposed baseline consists of a BEV encoder, a GFC as backbone,




Figure 4. Overall structure of the LLDN-GFC. There are five functions: (1), (2), (3), (4), and (5) indicate the BEV encoder, reshape & per-patch linear transform, Transformer or Mixer block, reshape & shared MLP, and detection head, respectively. H_{bev} , W_{bev} , C_{bev} , C_{out} , N_p^2 , and C_p , are the height, width, num. of channel of the pseudo-BEV image, num. of channel of the output pseudo-BEV image, the num. of total patches, and the num. of channel per patch in the global correlation, respectively.

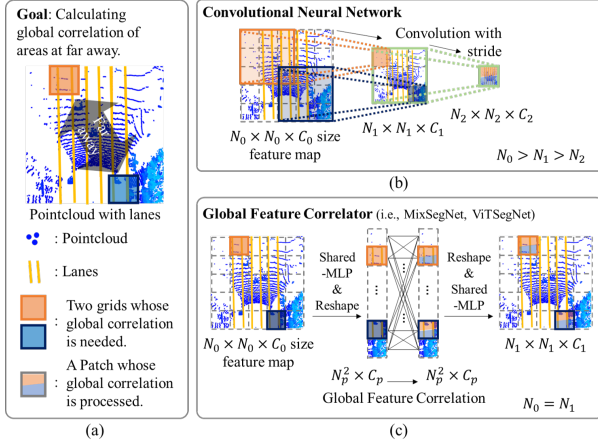


Figure 5. A comparison of global feature correlation between CNN and the proposed GFC, where N_0, N_1, N_2 and C_0, C_1, C_2 represent the size and the number of channels of the feature maps at three layers in depth order. (a) an example of two separated grids to calculate the global correlation, (b) the two grids contained in a feature map developed by CNN, and (c) correlation of the two grids in a feature map developed by GFC.

and a lane detection head, which are introduced in the following subsections.

BEV Encoder. The BEV encoder projects a 3D point cloud into a 2D pseudo-image and process it further to produce a 2D BEV feature map. We provide two variants of BEV encoder for the LLDN-GFC, namely point projector and pillar encoder.

The primary BEV encoder is the point projector [9, 21] that projects point clouds into the xy -horizontal plane and produces a BEV feature map using CNN. In order to maintain both high-resolution lane information and real-time speed, we design a ResNet-based CNN to output a feature map that is $1/8^2$ of the pseudo-image input.

An alternative for low computational 2D BEV encoder is the pillar encoder based on Point Pillars that has relatively small network size [10]. Pillar encoder has slightly lower performance but faster inference speed than the CNN-based point projector. Therefore, in this paper, pillar encoder is presented as an alternative for real-time applications. Details are in Appendix B.

GFC as Backbone. As shown in Fig. 5-a, lane lines on the road are thin, stretched along the entire point cloud, and only occupy a small number of pixels (i.e., sparse). Due to such thinness and sparsity, it is necessary to perform feature extraction in high resolution. In addition, the feature extractor should consider the correlation between distant grids within the BEV feature map. As such, we design our proposed GFC to calculate global correlations of the features in high resolution by utilizing patch-wise self-attention networks. We propose two variants of GFC: GFC-T (the main

proposal based on Transformer blocks [4]) and GFC-M (the low computational alternative based on Mixer blocks [24])

A major advantage of using patch-wise self attention networks is their capability to find correlations between distant grids (or patches) right from the early stages of backbone, as shown in Fig. 5-c. As such, the high-resolution information can be preserved (i.e., $N_0 = N_1$). This is in contrast with CNN-based feature extractors, which may find correlations between distant grids after several layers of convolutions and down samplings, thus lowering the resolution of information (i.e., $N_0 \gg N_2$), as shown in Fig. 5-b.

Quantitatively, we observe that patch-wise self-attention networks have higher performance compared to their CNN-based counterparts [14]. In addition, we visualize the qualitative results of intermediate feature maps and attention scores in Fig. 6 and 7, respectively. Both quantitative and qualitative results further indicate the aptness of using patch-wise self-attention networks for Lidar lane detection even on a relatively small number of data (i.e. 7687 training frames).

Detection Head and Loss Function. To design the detection head, we formulate the lane detection problem as a multi-class segmentation problem, where each pixel is assigned a class and a confidence score. The multi-class segmentation formulation enables the detection head to predict both lane classes and various lane shapes, which are important for motion planning where the ego vehicle need to plan inter-lane motions or recognize lane merging and separation. The LLDN-GFC detection head consists of two segmentation heads, each of which consists of a sequence of two-layer shared-MLP with a non-linear activation in-between.

As the number of lane samples are significantly smaller than the number of background samples on each frame, we incorporate the soft dice loss [22] for the confidence loss that inherently handles the imbalance problem. For the classification head, we choose the grid-wise cross-entropy loss [19] that has been widely used for multi-class classification problems, leading the network to learn to maximize the probability of the correct lane class during training. The total loss function is the summation of both the soft-dice loss and the cross-entropy loss as expressed in Appendix B.

4. Experiments and Comparison

In this section, we provide detailed performance comparisons between LLDN-GFC and conventional CNN-based LLDNs. In addition, we also discuss recent CLDNs for a general comparison to the LLDN-GFC performance.

Implementation Details. We evaluate two variants of LLDN-GFC, Proj28-GFC-T3 and Pillars-GFC-M5, which we observe during experiments (i.e., ablations in Appendix C) to have the best accuracy and speed-accuracy tradeoff, respectively. Proj28-GFC-T3 stands for LLDN-GFC with

Enc	Backbone	Total	Day	Night	Urban	Highway	Shp Curve	No Occ	Occ4~6	FPS
Proj-28	GFC-T3	82.1/81.1	82.2/81.4	82.0/80.7	81.7/80.6	82.5/81.7	78.0/76.7	82.9/81.9	75.9/75.5	11.6
	GFC-M3	79.7/78.8	79.9/79.0	79.6/78.4	78.9/77.7	80.8/80.0	74.6/72.6	80.4/79.4	72.5/71.7	13.3
	RNF-S13	73.2/70.5	72.6/70.1	74.0/71.0	73.1/70.4	73.3/70.6	70.5/68.1	74.9/72.3	63.5/59.0	13.1
	RNF-C13	78.0/75.3	77.6/75.1	78.5/75.5	77.7/74.8	78.3/76.01	76.0/73.1	79.6/77.0	69.3/65.3	13.0
	RNF-D23	72.1/68.8	71.3/68.3	73.0/69.4	71.9/68.7	72.3/69.0	69.6/66.5	74.0/70.7	61.9/57.6	12.7
Pillars	GFC-T5	78.5/77.3	78.5/77.6	78.4/77.0	77.8/76.4	79.2/78.4	72.5/70.2	79.4/78.2	70.2/69.5	13.8
	GFC-M5	74.8/73.5	74.8/73.6	74.9/73.4	72.0/70.5	78.2/77.1	64.6/62.2	75.5/74.2	65.2/62.3	16.3
	RNF-S13	64.6/18.2	62.9/16.4	66.5/20.4	59.4/15.6	70.7/21.4	51.1/13.1	65.5/19.2	44.9/4.7	15.7
	RNF-C13	76.8/40.6	75.9/39.1	77.8/42.4	74.5/40.6	79.6/40.6	67.6/32.5	77.9/43.6	62.5/20.4	15.5
	RNF-D23	63.2/19.0	61.6/17.0	65.0/21.2	58.6/17.2	68.6/21.1	51.1/13.5	64.2/20.2	43.9/4.9	15.2
Heuristic		26.4	26.8	26.0	23.3	29.6	27.6	28.1	16.7	18.3

Table 2. F1-score of confidence/classification for the proposed LLDN-GFC and various CNN-based LLDNs. Enc, Shp, Occ stands for BEV Encoder, sharp curve, and occlusion cases, respectively. We show no occlusion and severe occlusion (4~6 lanes occluded) cases, while other occlusion levels are presented in Appendix C. FPS stands for frame per second, which represents the overall computational cost (FLOPs, data efficiency, etc.) of the networks during inference, similar to throughput in [24]. Note that we only show F1-score of confidence for the heuristic technique.

point projector encoder with 28 layers and GFC with three Transformer blocks. Pillars-GFC-M5 stands for LLDN-GFC with pillar encoder and GFC with five Mixer blocks.

We use RTX3090 GPUs for training the networks on the K-Lane for 60 epochs using Adam optimizer [8] with a batch size 4 and a learning rate of 0.0002. All training and evaluations are implemented with PyTorch 1.7.1 [17] on an Ubuntu 18.04 machine.

4.1. LLDN-GFC vs. CNN-based LLDN

We consider three types of CNN-based backbone to be compared with the proposed GFC, namely: RNF-S13, RNF-D23, and RNF-C13, where (1) RNF represent ResNet [6] with feature pyramid network (FPN) [11], (2) S13, D23, and C13 represent residual blocks implemented with strided convolution, dilated convolution, and CBAM [28] of 13 or 23 layers, respectively. The model capacities of the counterparts are also determined with experiments (i.e., ablations in Appendix C). The FPN is applied to synthetically consider feature maps of different levels, and the dilated convolution increase the receptive field without loss of resolution, which is utilized in the existing LLDN [14]. The CBAM performs self-attention mechanism similar to the proposed LLDN-GFC, but applied with per-channel convolution operation, meaning that it does not perform global correlations for all patches as in LLDN-GFC. For this reason, LLDN with RNF-C shows lower performance than the proposed LLDN-GFC, as shown in Table 2.

As summarized in Table 2, the proposed LLDN-GFC shows superior performance than the LLDNs with conventional CNN-based backbone of various depths. In particular, LLDN-GFC shows robust performance against sever occlusions, where four or more lanes are occluded.

Fig. 6 show qualitative assessment of the robustness of LLDN-GFC based on the visualization of intermediate

feature maps. We can observe on the heatmaps that both Proj28-GFC-T3 (a) and Proj28-GFC-M3 (b) clearly extract lanes with better resolution, especially on the deeper layers. This is in contrast with CNN-based LLDN, shown in (c) and (d), where the lanes tend to blur. In other words, the lane features extracted by Proj28-GFC-T3 and Proj28-GFC-M3 are more distinctive to the backgrounds compared to the lane features extracted by CNN-based LLDNs. In

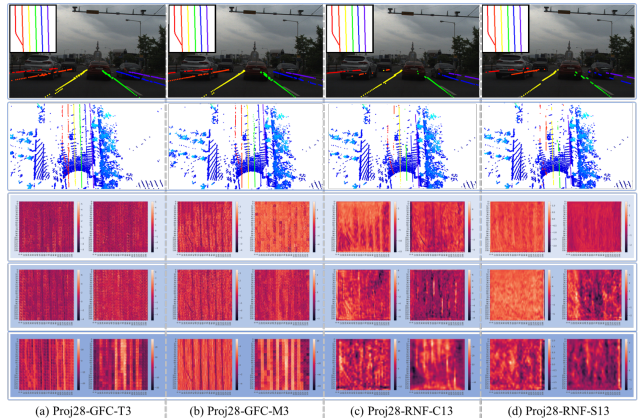


Figure 6. Comparison of the lane detection performance between the LLDN-GFC and CNN-based LLDNs for occluded lanes condition. The four columns are inference results of (a) Proj28-GFC-T3, (b) Proj28-GFC-M3, (c) Proj28-RNF-C13, and (d) Proj28-RNF-S13. The first row shows the projection of inference results onto the front view image with labels in the upper left corner, while the second row shows the inference on the BEV point cloud. From the third to fifth row, we show the heatmaps sampled along the channels of the 1st, 2nd, and 3rd block output feature map of the backbones (e.g., 1st, 2nd, and 3rd Transformer blocks of GFC or residual blocks of RNF). Heatmaps for various scenarios such as curved and merging lane lines are introduced in Appendix D.

addition, even in the presence of occlusions, Proj28-GFC-T3 and Proj28-GFC-M3 are capable of predicting the lane shapes through correlations with the non-occluded lanes, which are not observed in the CNN-based LLDNs.

4.2. LLDN-GFC Attention Visualization

In this subsection, we discuss the robustness of the proposed LLDN-GFC, Proj28-GFC-T3, against the occlusion scenario using the visualization of attention score.

The proposed GFC is based on the self-attention mechanism that utilizes correlations between data units to make the network give more attention to the meaningful region on the feature map. As such, we can see the region that is considered as important by the GFC-T3 by visualizing the attention score produced by each Transformer blocks, as shown in Fig. 7.

From the visualization, we can see that the network give more attention to the regions which contain lane lines by attenuating the magnitude of non-lane-lines (irrelevant) features. As the layers get deeper, the network expands its region of interests, indicated by the increasing area with high attention scores. We observe that utilizing three blocks of transformer for GFC-T3 is sufficient to ensure the self-attention mechanism to cover the entire region of the point cloud which contains lane lines. Additionally, note that for the query location (the yellow box in Fig 7), the network produce high attention scores to regions in which lane lines are present, even if the query location is occluded. Such phenomena indicates the robustness of LLDN-GFC to occlusions, where predictions are made by considering the entire point cloud such that occluded lane lines can still be estimated accurately. This may not be possible for CNN-based LLDN, which features are recognized through local convolutions.

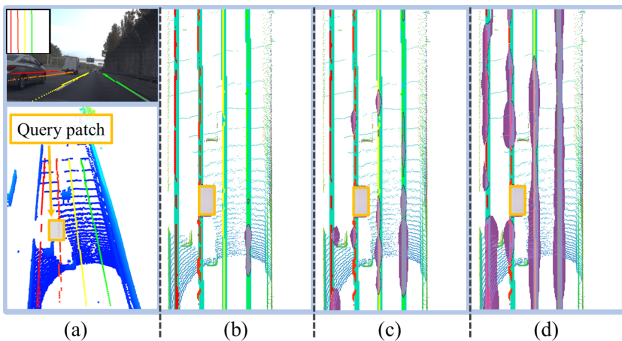


Figure 7. Attention score visualization of Proj28-GFC-T3. (a-upper) shows the projection of inference results onto the front view image with labels in the upper left corner, while (a-lower) shows the inference result and the current query patch (yellow box) on top of the BEV point cloud. (b) to (d) show the point cloud in BEV, lane inference results, query patch, labels in cyan, and attention score in purple for block 1, 2, and 3 of the GFC, respectively.

4.3. LLDN-GFC vs. Camera Lane Detection

Most state-of-the-art lane detection networks in the literature are for front-view camera images. This means most CLDNs are trained to detect lane lines in the front-view map. On the other hand, LLDN-GFC is trained to detect lane lines in the BEV map. In addition, the environment in which the data are collected is different. CULane is composed of data mostly for urban roads, while K-Lane consists of data for both urban roads and highways. Since these Lidar and camera datasets do not use the same representations and are collected in different environments, we cannot compare the CLDNs simply using the reported performance in the literature.

However, recent CLDNs show a significant performance drop for the night time data comparing to the daytime data. For example, CondLaneNet-Large [13], LaneATT-Large [23], and CurveLane-NAS-L [29] show 18.67%, 20.93%, and 21.8% drops between daytime and nighttime conditions, respectively. In contrast, as shown in Table 1, the proposed LLDN-GFC shows almost no performance degradation (only 0.2% difference). This is because the Lidar is robust to light conditions, which demonstrates that LLDN is a reliable function for autonomous driving.

5. Conclusion

In this work, we introduce K-Lane which, to the best of our knowledge, is the world’s first publicly available dataset for Lidar lane detection. K-Lane consists of over 15K high-quality annotated Lidar data in diverse and challenging driving conditions, along with well-calibrated front-view RGB images. The driving conditions include various lighting (daytime and nighttime), lane occlusions (up to 6 occluded lane lines), and road structures (merging, diverging, curved lanes). In addition, we provide the development kits for K-Lane including the annotation, visualization, training, and benchmarking tools. We also introduce a baseline network for Lidar lane detection, which we term LLDN-GFC. LLDN-GFC utilizes self-attention mechanisms to extract lane features via global correlation, and show superior performance compared to the conventional CNN-based LLDNs. In addition, we show the importance of Lidar lane detection networks, where there is only little performance degradation in between detection in the daytime and detection in the nighttime, in contrast to camera-based lane detection networks. As such, we expect this work to pave the way for further studies in the field of Lidar lane detection, and improve the safety aspects in autonomous driving.

Acknowledgment

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2021R1A2C3008370)

References

- [1] Min Bai, Gellert Mattyus, Namdar Homayounfar, Shenlong Wang, Shrinidhi Kowshika Lakshmikanth, and Raquel Urtasun. Deep multi-sensor lane detection. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3102–3109. IEEE, 2018. 1, 2, 3, 5
- [2] Ganlu Deng and Yefu Wu. Double lane line edge detection method based on constraint conditions hough transform. In *2018 17th International Symposium on Distributed Computing and Applications for Business Engineering and Science (DCABES)*, pages 107–110. IEEE, 2018. 1, 3
- [3] Yue Dong, Jintao Xiong, Liangchao Li, and Jianyu Yang. Robust lane detection and tracking for lane departure warning. In *2012 International Conference on Computational Problem-Solving (ICCP)*, pages 461–464. IEEE, 2012. 1, 3
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. 2, 4, 6
- [5] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996. 3
- [6] 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. 3, 7
- [7] Danilo Caceres Hernandez, Van-Dung Hoang, and Kang-Hyun Jo. Lane surface identification based on reflectance using laser range finder. In *2014 IEEE/SICE International Symposium on System Integration*, pages 621–625. IEEE, 2014. 3
- [8] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 7
- [9] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–8. IEEE, 2018. 6
- [10] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. 6
- [11] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 7
- [12] Philipp Lindner, Eric Richter, Gerd Wanielik, Kiyokazu Takagi, and Akira Isogai. Multi-channel lidar processing for lane detection and estimation. In *2009 12th International IEEE Conference on Intelligent Transportation Systems*, pages 1–6. IEEE, 2009. 3
- [13] Lizhe Liu, Xiaohao Chen, Siyu Zhu, and Ping Tan. Condlanenet: a top-to-down lane detection framework based on conditional convolution, 2021. 1, 3, 8
- [14] Philipp Martinek, Gheorghe Pucea, Qing Rao, and Udharaj Sivalingam. Lidar-based deep neural network for reference lane generation. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 89–94. IEEE, 2020. 2, 3, 5, 6, 7
- [15] Ouster. Os2-64. <https://ouster.com/products/os2-lidar-sensor/>, 2020. Accessed: 2021-09-08. 4
- [16] Xingang Pan, Jianping Shi, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Spatial as deep: Spatial cnn for traffic scene understanding. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 1, 3, 4, 5
- [17] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. 7
- [18] Zequn Qin, Huanyu Wang, and Xi Li. Ultra fast structure-aware deep lane detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 276–291. Springer, 2020. 3
- [19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 6
- [20] David Schreiber, Bram Alefs, and Markus Clabian. Single camera lane detection and tracking. In *Proceedings. 2005 IEEE Intelligent Transportation Systems, 2005.*, pages 302–307. IEEE, 2005. 1
- [21] Martin Simony, Stefan Milzy, Karl Amendey, and Horst-Michael Gross. Complex-yolo: An euler-region-proposal for real-time 3d object detection on point clouds. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 6
- [22] Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. *Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations*. Springer, 2017. 6
- [23] Lucas Tabelini, Rodrigo Berriel, Thiago M Paixao, Claudine Badue, Alberto F De Souza, and Thiago Oliveira-Santos. Keep your eyes on the lane: Real-time attention-guided lane detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 294–302, 2021. 1, 8
- [24] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision, 2021. 2, 4, 6, 7
- [25] TuSimple. Tusimple benchmark. <https://github.com/TuSimple/tusimple-benchmark>, 2017. Accessed: 2021-09-08. 1, 3, 5
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia

- Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 4
- [27] Chunxiang Wang, Hengrun Zhang, Ming Yang, Xudong Wang, Lei Ye, and Chunzhao Guo. Automatic parking based on a bird’s eye view vision system. *Advances in Mechanical Engineering*, 6:847406, 2014. 1
- [28] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 3, 7
- [29] Hang Xu, Shaoju Wang, Xinyue Cai, Wei Zhang, Xiaodan Liang, and Zhenguo Li. Curvelane-nas: Unifying lane-sensitive architecture search and adaptive point blending. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pages 689–704. Springer, 2020. 8
- [30] Brandon Yang, Gabriel Bender, Quoc V. Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché Buc, Edward A. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 1305–1316, 2019. 3