

# ONCE-3DLanes: Building Monocular 3D Lane Detection

Fan Yan<sup>1\*</sup> Ming Nie<sup>1\*</sup> Xinyue Cai<sup>2</sup> Jianhua Han<sup>2</sup> Hang Xu<sup>2</sup> Zhen Yang<sup>2</sup>  
Chaoqiang Ye<sup>2</sup> Yanwei Fu<sup>1</sup> Michael Bi Mi<sup>2</sup> Li Zhang<sup>1†</sup>

<sup>1</sup>School of Data Science, Fudan University <sup>2</sup>Huawei Noah's Ark Lab

<https://once-3dlanes.github.io>

## Abstract

We present *ONCE-3DLanes*, a real-world autonomous driving dataset with lane layout annotation in 3D space. Conventional 2D lane detection from a monocular image yields poor performance of following planning and control tasks in autonomous driving due to the case of uneven road. Predicting the 3D lane layout is thus necessary and enables effective and safe driving. However, existing 3D lane detection datasets are either unpublished or synthesized from a simulated environment, severely hampering the development of this field. In this paper, we take steps towards addressing these issues. By exploiting the explicit relationship between point clouds and image pixels, a dataset annotation pipeline is designed to automatically generate high-quality 3D lane locations from 2D lane annotations in 211K road scenes. In addition, we present an extrinsic-free, anchor-free method, called SALAD, regressing the 3D coordinates of lanes in image view without converting the feature map into the bird's-eye view (BEV). To facilitate future research on 3D lane detection, we benchmark the dataset and provide a novel evaluation metric, performing extensive experiments of both existing approaches and our proposed method.

The aim of our work is to revive the interest of 3D lane detection in a real-world scenario. We believe our work can lead to the expected and unexpected innovations in both academia and industry.

## 1. Introduction

The perception of lane structure is one of the most fundamental and safety-critical tasks in autonomous driving system. It is developed with the desired purpose of preventing accidents, reducing emissions and improving the traffic efficiency [4]. It serves as a key roll of many applications, such

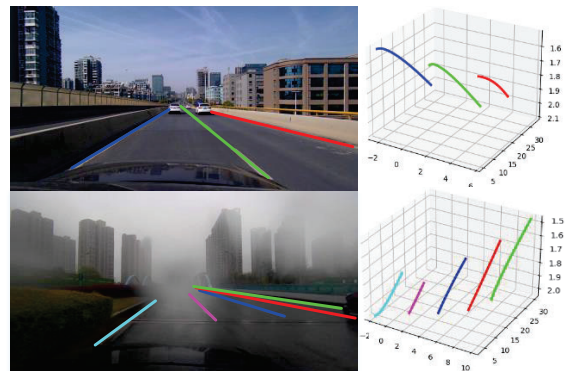


Figure 1. Images and 3D lane examples of *ONCE-3DLanes* dataset. *ONCE-3DLanes* covers various locations, illumination conditions, weather conditions and with numerous slope scenes.

as lane keeping, high-definition (HD) map modeling [13], trajectory planning *etc.* In light of its importance, there has been a recent surge of interest in monocular 3D lane detection [6, 7, 10, 15, 32]. However, existing 3D lane detection datasets are either unpublished, or synthesized in a simulated environment due to the difficulty of data acquisition and high labor costs of annotation. With the only synthesized data, the model inevitably lacks generalization ability in real-world scenarios. Although benefiting from the development of domain adaptation method [8], it still cannot completely alleviate the domain gap.

Most existing image-based lane detection methods have exclusively focused on formulating the lane detection problem as a 2D task [29, 34, 35], in which a typical pipeline is to firstly detect lanes in the image plane based on semantic segmentation or coordinate regression and then project the detected lanes in top view by assuming the ground is flat [28, 33]. With well-calibrated camera extrinsics, the inverse perspective mapping (IPM) is able to obtain an acceptable approximation for 3D lane in the flat ground plane. However, in real-world driving environment, roads are not always flat [7] and camera extrinsics are sensitive to vehicle body motion due to speed change or bumpy road, which will

\*The first two authors contributed equally to this work.

†Li Zhang (lizhangfd@fudan.edu.cn) is the corresponding author at School of Data Science and Shanghai Key Laboratory of Intelligent Information Processing, Fudan University

lead to the incorrect perception of the 3D road structure and thus unexpected behavior may happen to the autonomous driving vehicle.

To overcome above shortcomings associated with flat-ground assumption, 3D-LaneNet [7], directly predicts the 3D lane coordinates in an end-to-end manner, in which the camera extrinsics are predicted in a supervised manner to facilitate the projection from image view to top view. In addition, an anchor-based lane prediction head is proposed to produce the final 3D lane coordinates from the virtual top view. Despite the promising result exhibits the feasibility of this task, the virtual IPM projection is difficult to learn without the hard-to-get extrinsics and the model is trained under the assumption that zero degree of camera roll towards the ground plane. Once the assumption is challenged or the need of extrinsic parameters is not satisfied, this method can barely work.

In this work, we take steps towards addressing above issues. For the first time, we present a real-world 3D lane detection dataset *ONCE-3DLanes*, consisting of 211K images with labeled 3D lane points. Compared with previous 3D lane datasets, our dataset is the largest real-world lane detection dataset published up to now, containing more complex road scenarios with various weather conditions, different lighting conditions as well as a variety of geographical locations. An automatic data annotation pipeline is designed to minimize the manual labeling effort. Comparing to the method [7] of using multi-sensor and expensive HD maps, ours is simpler and easier to be implemented. In addition, we introduce a spatial-aware lane detection method, dubbed *SALAD*, in an extrinsic-free and end-to-end manner. Given a monocular input image, *SALAD* directly predicts the 2D lane segmentation results and the spatial contextual information to reconstruct the 3D lanes without explicit or implicit IPM projection.

The contributions of this work are summarized as follows: (i) For the first time, we present a largest 3D lane detection dataset *ONCE-3DLanes*, alongside a more generalized evaluation metric to revive the interest of such task in a real-world scenario; (ii) We propose a method, *SALAD* that directly produce 3D lane layout from a monocular image without explicit or implicit IPM projection.

## 2. Related work

### 2.1. 2D lane detection

There are various methods [1, 14, 22, 29, 31, 35] proposed to tackle the problem of 2D lane detection. Segmentation-based methods [18, 24, 28, 29] predict pixel-wise segmentation labels and then cluster the pixels belonging to the same label together to predict lane instances. Proposal-based methods [34, 35] first generate lane proposals either from the vanishing points [33] or from the edges of image [21],

and then optimize the lane shape by regressing the lane offset. There are some other methods [12, 28, 33, 37] trying to project the image into the top view and using the properties that lanes are almost parallel and can be fitted by lower order polynomial in the top view to fit lanes. However, most methods are limited in the image view, lack the image-to-world step or suffer from the untenable flat-ground assumption. As a result, formulating lane detection as a 2D task may cause inappropriate behaviors for autonomous vehicle when encountering hilly or slope roads [15].

### 2.2. 3D lane detection

**LiDAR-based lane detection.** Several methods [16, 17, 36] have been proposed using LiDAR to detect 3D lanes, [11] use the characteristic that the intensity values of different material are different to filter out the point clouds of the lanes by a certain intensity threshold and then cluster them to obtain 3D lanes. However, it's hard to determine the specific intensity threshold since the material used in different countries or regions is distinct, and the intensity value varies much in various weather conditions, *e.g.*, rainy or snowy.

**Multi-sensor lane detection.** Other methods [2, 39] try to aggregate information from both camera and LiDAR sensors to tackle lane detection task. Specifically, [26] predicts the ground height from LiDAR points to project the image to the dense ground. It combines the image information with LiDAR information to produce lane boundary detection results. Nevertheless, it's difficult to guarantee that the image and the point clouds appear in pairs in the real scenes, *e.g.*, CULane dataset only contains images.

**Monocular lane detection.** Recently there are a few methods [6, 7, 10, 15, 32] trying to address this problem by directly predicting from a single monocular image. The pioneering work 3D LaneNet [7] predicts camera extrinsics in a supervised manner to learn the inverse perspective mapping(IPM) projection, by combining the image-view features with the top-view features. Gen-LaneNet [10] proposed a new geometry-guided lane anchor in the virtual top view. By decoupling the learning of image segmentation and 3D lane prediction, it achieves higher performance and are more generalizable to unobserved scenes. Instead of associating each lane with a predefined anchor, 3D-LaneNet+ [32] proposes an anchor-free, semi-local representation method to represent lanes. Although the ability to detect more lane topology structures shows the anchor-free method's power. However, all the above methods need to learn a projection matrix in a supervised way to align the image-view features with top-view features, which may cause height information loss. While our proposed method directly regresses the 3D coordinates in the image view without considering camera extrinsics.

Datasets	images real published		weather geographical	
			change	locations
synthetic-3D-lanes	306K	No	Yes	No
Apollo-Sim-3D	10.5K	No	Yes	No
Real-3D-lanes	85K	Yes	No	-
Ours	211K	Yes	Yes	Yes
				highway,urban residential area highway,rural highway,bridge tunnel,suburb downtown

Table 1. Comparison of different 3D lane detection datasets. ”-” means not mentioned. Ours is the first published real-world dataset covering different weather conditions and geographical locations.

### 2.3. Lane datasets

Existing 3D lane detection datasets are either unpublished or synthesized in simulated environment. GenLanenet [10] uses Unity game engine to build 3D worlds and releases a synthetic 3D lane dataset, Apollo-Sim-3D, containing 10.5K images. 3D-LaneNet [7] adopts a graphics engine to model terrains using a Mixture of Gaussians distribution. Lanes modeled by a 4<sup>th</sup> degree polynomial in top view are placed on the terrains to generate the synthetic 3D lanes dataset *synthetic-3D-lanes*, containing 306K images with a resolution of 360 × 480. A real 3D lane dataset *Real-3D-lanes* with 85K images is also created using multiple sensors including camera, LiDAR scanner and IMU as well as the expensive HD maps in [7].

In this paper, we publish the first real-world 3D lane dataset *ONCE-3DLanes* which contains 211K images and covers abundant scenes with various weather conditions, different lighting conditions as well as a variety of geographical locations. Comprehensive comparisons of 3D lane detection datasets are shown in Table 1.

## 3. ONCE-3DLanes

### 3.1. Dataset introduction

**Raw data.** We construct our *ONCE-3DLanes* dataset based on the most recent large-scale autonomous driving dataset ONCE (one million scenes) [25] considering its superior data quality and diversity. ONCE contains 1 million scenes and 7 million corresponding images, the 3D scenes are recorded with 144 driving hours covering different time periods including morning, noon, afternoon and night, various weather conditions including sunny, cloudy and rainy days, as well as a variety of regions including downtown, suburbs, highway, bridges and tunnels. Since the camera data is captured at a speed of two frames per second and most adjacent frames are very similar, we take one frame every five frames to build our dataset to reduce data redundancy. Also, the distortions are removed to enhance the image quality and improve the projection accuracy from LiDAR to camera.

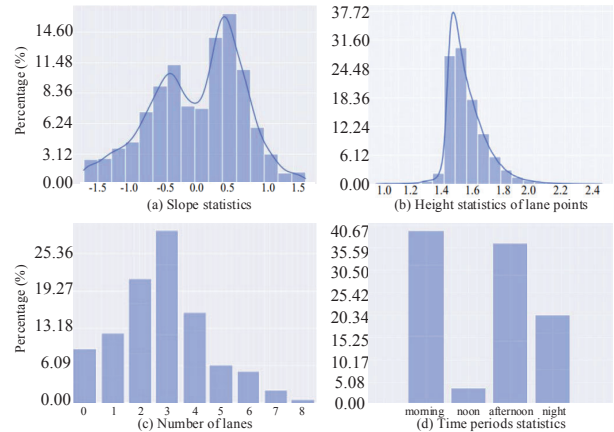


Figure 2. An overview of slope scenes statistics is shown in (a). The distribution of the height of lane points is shown in (b). The histogram of the average number of lanes per image and the time periods statistics are shown in (c) and (d) respectively.

Thus, by downsampling the ONCE five times, our dataset contains 211k images taken by a front-facing camera.

**Lane representation.** A lane  $L^k$  in 3D space is represented by a series of points  $\{(x_i^k, y_i^k, z_i^k)\}_{i=1}^n$ , which are recorded in the 3D camera coordinate system with unit meter. The camera coordinate system is placed at the optical center of the camera, with X-axis positive to the right, Y-axis downward and Z-axis forward.

**Dataset analysis.** The projection error from front-view to top-view mainly occurs in the situation of slope ground, so we focus on analyzing the slope statistics on *ONCE-3DLanes*. The mean slope of lanes in each scene is utilized to represent the slope of this scene. The slope of a specific lane in forward direction which is considered to be the most important is calculated as follows:

$$slope = (y_2 - y_1)/(z_2 - z_1) \quad (1)$$

where  $(x_1, y_1, z_1)$  and  $(x_2, y_2, z_2)$  are the start point and the end point of the lane respectively. The distribution of the slope conditions and the histogram of the number of lanes per image are shown in Figure 2. It shows that our dataset is full of complexity and contains enough various slope scenes with different illumination conditions.

**Dataset splits.** Follow the ONCE dataset, our benchmark contains the same 3K scenes for validation and 8K scenes for testing. To fully make use of the raw data, the training dataset not only contains the original 5K scenes, but also the unlabeled 200K scenes.

### 3.2. Annotation pipeline

Lanes are a series of points on the ground, which are hard to be identified in point clouds. Hence the high-quality

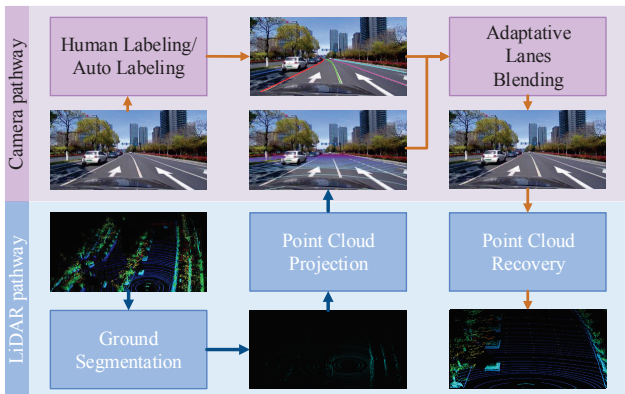


Figure 3. **Dataset Annotation Pipeline:** With the paired image and Lidar point clouds as input, the 2D lanes on the image are firstly labeled and broadened to get the lane regions; secondly the ground points in the point clouds are filtered out through ground segmentation; thirdly the filtered ground points are projected to the image and collect the points which are contained in the lane regions, finally cluster the points to get the real lane points.

annotations of 3D lanes are expensive to obtain, whilst it is much cheaper to annotate the lanes in 2D images. The paired LiDAR point clouds and image pixels are thoroughly investigated and used to construct our 3D Lane dataset. An overview of dataset construction pipeline is shown in Figure 3. The pipeline consists of five steps: Ground segmentation, Point cloud projection, Human labeling/Auto labeling, Adaptive lanes blending and Point cloud recovery. These steps are described in detail below.

**Ground segmentation.** Lanes are painted on the ground, which is a strong prior to locate the precise coordinate in the 3D space. To make full use of human prior and avoid the reflection of point clouds aliasing between lanes and other objects, the ground segmentation algorithm is utilized to get the ground LiDAR points at first.

The ground segmentation is performed in a coarse-to-fine manner. In the coarse way, since the height of the LiDAR points reflected by the ground always settles in certain intervals, a pre-defined threshold is adopted to filter out those points lying on the ground coarsely based on the height statistics of the LiDAR points among the whole dataset as seen in Figure 2(b). In the fine way, several points in front of the vehicle are sampled randomly as seeds and then the classic region growth method is applied to get the fine segmentation result.

**Point cloud projection.** In this step, the previous extracted ground LiDAR points are projected to the image plane with the help of calibrated LiDAR-to-camera extrinsics and camera intrinsics based on the classic homogeneous transformation, which reveals the explicit corresponding relationship between the 3D ground LiDAR points and the 2D ground pixels in the image.

**Human labeling / Auto labeling.** To obtain 2D lane labels in the images and alleviate the taggers’ burden, a robust 2D lane detector which is trained in million-level scenes is firstly used to automatically pre-annotate pseudo lane labels. At the same time, professional taggers are required to verify and correct the pseudo labels to ensure the annotation accuracy and quality.

**Adaptative lanes blending.** After getting the accurate 2D lane labels and ground points, in order to judge whether a ground point belongs to the lane marker or not, we broaden 2D lane labels with an appropriate and adaptive width to get lane regions. Due to the perspective principle, a lane is broadened with different widths according to the distance from camera. After the point clouds projection procedure, we consider the ground points which are contained in the lane regions as the lane point clouds.

**Point cloud recovery.** Finally, we select these lane point clouds out. And for a specific lane, the lane point clouds in the same beam are clustered to get the lane center points to represent this lane.

To ensure the accuracy of annotations, we do not interpolate between center points in data collection stage. While in training stage, we use cubic spline interpolation to generate dense supervised labels. We also compared our annotations with manually labeling results on a small portion of data, which shows the high quality of our annotations. The interpolation code will be made public along with our dataset.

## 4. SALAD

In this section, we introduce *SALAD*, a spatial-aware monocular lane detection method to perform 3D lane detection directly on monocular images. In contrast to previous 3D lane detection algorithms [7, 10, 32], which project the image to top view and adopt a set of predefined anchors to regress 3D coordinates, our method does not require human-crafting anchors and the supervision of extrinsic parameters. Inspired by SMOKE [23], *SALAD* consists of two branches: semantic awareness branch and spatial contextual branch. The overall structure of our model is illustrated in Figure 4. In addition, we also adopt a revised joint 3D lane augmentation strategy to improve the generalization ability. The details of our network architecture and augmentation methods are discussed in the following parts.

### 4.1. Backbone

We choose the Segformer [38] as our backbone to extract the global contextual features and to learn the slender structure of lanes. To be concrete, given an image  $I \in \mathbb{R}^{H \times W \times 3}$ , the hierarchical transformer backbone encodes the image  $I$  into multi-level features at  $\{1/4, 1/8, 1/16, 1/32\}$  of the input resolution. Then all multi-level features are upsampled to  $\frac{H}{4} \times \frac{W}{4}$  and concatenated through a MLP-based decoder and a convolutional feature fusion layer to aggregate

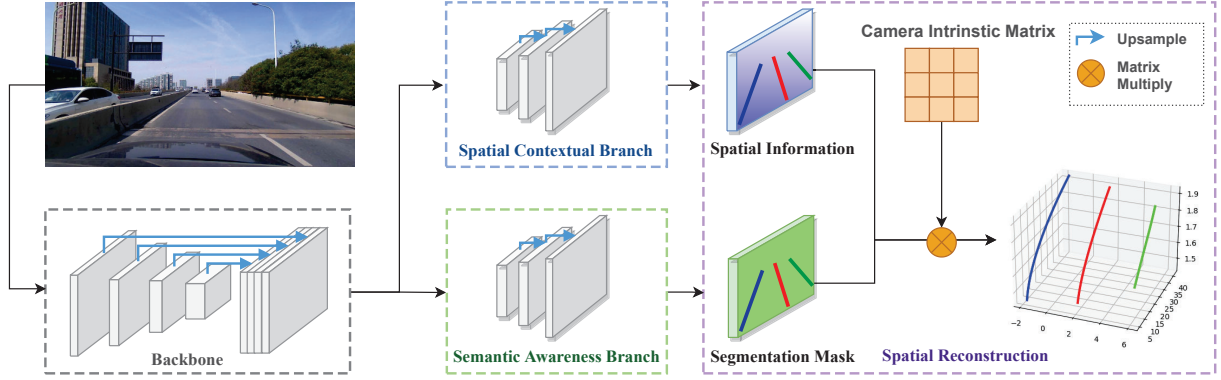


Figure 4. **The architecture of SALAD.** The backbone encodes an input image into deep features and two branches namely semantic awareness branch and spatial contextual branch decode the features to get the lane’s spatial information and the segmentation mask. Then 3D reconstruction is performed by integrating these information and finally obtain the 3D lane positions in real-world scene.

the multi-level features into  $F \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$ . Specifically, we adopt Segformer-B2 as our feature extractor.

#### 4.2. Semantic awareness branch

Traditional 3D lane detection methods [7, 10, 32] directly projecting feature map from front view to top view, are not reasonable and harmful to the performance of prediction as the feature map may not be organized following the perspective principle.

In order to directly regress 3D lane coordinates, we first design a semantic awareness branch to make full use of 2D semantic information, which provides 2D lane point proposals to aggregate 3D information together. It is a relatively easy task to extract 2D features of lane markers from images with rich semantic information. In addition, current mature experience on semantic segmentation task can also be used to enhance our semantic awareness branch. We follow the method of [28], to encode 2D lane coordinates  $\{(u_i^k, v_i^k)\}_{i=1}^n$  into ground-truth segmentation map  $S_{gt} \in \mathbb{R}^{H \times W \times 1}$ .

During training time, image  $I \in \mathbb{R}^{H \times W \times 3}$  and ground-truth  $S_{gt} \in \mathbb{R}^{H \times W \times 1}$  pairs are utilized to train semantic awareness branch. During inference, given an image  $I \in \mathbb{R}^{H \times W \times 3}$ , we are able to locate the foreground lane points on the binary mask  $S \in \mathbb{R}^{H \times W \times 1}$ .

According to [5], inverse projection from 2D image to 3D space is an underdetermined problem. Based on the segmentation map generated by semantic awareness branch, spatial information of each pixel is also required to transfer this segmentation map from 2D image plane to 3D space.

#### 4.3. Spatial contextual branch

To reconstruct 3D lanes from 2D lane points generated by semantic awareness branch, we propose the spatial contextual branch to predict vital 3D offsets. To sum up, our

spatial contextual branch predicts a regression result  $O = [\delta_u, \delta_v, \delta_z]^T \in \mathbb{R}^{3 \times H \times W}$ .  $\delta_u$  and  $\delta_v$  denote the pixel location offsets of lane points predicted in segmentation branch ( $u_s, v_s$ ), and generate accurate 2D lane positions. The  $\delta_z$  denotes the pixel-level prediction of depth.

Due to the downsampling and lacking of global information, the locations of the predicted lane points are not accurate enough. Our spatial contextual branch accepts feature  $F$  and outputs a pixel-level offset map, which predicts the spatial location shift  $\delta_u$  and  $\delta_v$  of lane points along  $u$  and  $v$  axis on the image plane. With the predictions of pixel location offsets  $\delta_u$  and  $\delta_v$ , the rough estimation of lane point locations is modified by global spatial context:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} u_s + \delta_u \\ v_s + \delta_v \end{bmatrix}. \quad (2)$$

In order to recover 3D lane information, the spatial contextual branch also generates a dense depth map to regress on depth offset  $\delta_z$  for each pixel of the lane markers. Considering the depth of the ground on image plane increases along rows, we assign each row of the depth map a predefined shift  $\alpha_r$  and scale  $\beta_r$ , and perform regression in a residual way. The standard depth value  $z$  is recovered as following:

$$z = \alpha_r + \beta_r \delta_z. \quad (3)$$

The ground-truth depth map is generated by projecting 3D lane points  $\{(x_i^k, y_i^k, z_i^k)\}_{i=1}^n$  on the image plane to get pixel coordinates  $\{(u_i^k, v_i^k, z_i^k)\}_{i=1}^n$ . Then at each pixel  $(u_i^k, v_i^k)$ , its corresponding depth value is assigned to  $z_i^k$ . Following [20], we apply depth completion on the sparse depth map to get the dense depth map  $D_{gt}$  to provide sufficient training signals for our spatial contextual branch.

#### 4.4. Spatial reconstruction

The spatial information predicted by our spatial contextual branch of our model plays a virtual role in 3D lane reconstruction. To map 2D lane coordinates back to 3D spatial location in camera coordinate system, the depth information is an indispensable element. To be concrete, given the camera intrinsic matrix  $K_{3 \times 3}$ , a 3D point  $(x, y, z)$  in camera coordinate system can be projected to a 2D image pixel  $(u, v)$  as:

$$z \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = K_{3 \times 3} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{pmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix}, \quad (4)$$

where  $f_x$  and  $f_y$  represent the focal length of the camera,  $(c_x, c_y)$  is the principal point and  $s$  is the axis skew. Thus, given a 2D lane point in the image with pixel coordinates  $(u, v)$  along with its depth information  $d$ , noted that the depth denotes the distance to the camera plane, so the depth  $d$  is the same as the  $z$  in the camera coordinate system. Thus 3D lane point in camera coordinate system  $(x, y, z)$  can be restored as follows:

$$\begin{cases} z = d, \\ y = \frac{z}{f_y} \cdot (v - c_y), \\ x = \frac{z}{f_x} \cdot [(u - c_x) - \frac{s}{f_y} (v - c_y)]. \end{cases} \quad (5)$$

Utilizing the fixed parameters of the camera intrinsic, we can project the 2D lane proposal points back to 3D locations to reconstruct our 3D lanes.

#### 4.5. Loss function

Given an image and its corresponding ground-truth 3D lanes, the loss function between predicted lanes and ground-truth lanes are formulated as:

$$\mathcal{L} = \mathcal{L}_{seg} + \lambda \mathcal{L}_{reg}. \quad (6)$$

$\mathcal{L}_{seg}$  is for the binary segmentation branch, with a cross-entropy loss in a pixel-wise manner on the segmentation map. For a specific pixel in segmentation map  $S$ ,  $y_i$  is the label and  $p_i$  is the probability of being foreground pixels:

$$\mathcal{L}_{seg} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)]. \quad (7)$$

$\mathcal{L}_{reg}$  is for the spatial contextual branch, which predicts the spatial offsets  $O = [\delta_u, \delta_v, \delta_z]^T$ . we choose smooth L1 loss to regress these spatial contextual information  $O$ :

$$\mathcal{L}_{reg} = \frac{1}{N} \sum_{i=1}^N [smooth_{L1}(\hat{O}_i - O_i)]. \quad (8)$$

$\lambda$  denotes the penalty term for regression loss and is set to 1 in our experiments.

#### 4.6. Data augmentation

Randomly horizontal flip and image scaling are common data augmentation methods to improve the generalization ability of 2D lane detection models. However, it is worth noting that the image shift and scale augmentation methods will cause 3D information inconsistent with the data augmentation [23]. We revise it by proposing a joint scale strategy and the augmentation steps are as follows:

Given an image with height  $H$  and weight  $W$ , the image center is  $(c_u, c_v) = (\frac{W}{2}, \frac{H}{2})$ . With a fixed camera intrinsic matrix, the principal center is  $(c_x, c_y)$  as mentioned before. To guarantee the 3D point to align with the pixel when scaling, it is proved that the scaling center  $(s_x, s_y)$  is

$$\begin{cases} s_x = c_x + s \cdot (c_u - c_x), \\ s_y = c_y + s \cdot (c_v - c_y). \end{cases} \quad (9)$$

To ensure we can restore the same size of the original image from the scaled image, we first crop the top of the image with size  $c$ . Then we scale the cropped image according to the scaling center with the proportion  $s$ . The relationship of 3D information of a specific pixel before scaling  $(x, y, z)$  and after scaling  $(\hat{x}, \hat{y}, \hat{z})$  is:

$$(\hat{x}, \hat{y}, \hat{z}) = (x, y, z \cdot s). \quad (10)$$

Take scaling factor  $s$  which is less than one for example. If the image is scaled with factor  $s$ , in the camera coordinate system, it is like the camera is moving forward in the  $Z$  direction. As for a specific point, the  $x$  and  $y$  keep the same while the  $z$  become smaller and the new  $z$  equals  $z \cdot s$ .

### 5. Experiment

In this section, our experiments are presented as follows. First we introduce our experimental setups, including evaluation metrics and implementation details. Then we evaluate our baseline method on our *ONCE-3DLanes* dataset and investigate the evaluation performance of different hyperparameter settings. Next we compare our proposed method with the prior state-of-the-art to prove the superiority of our proposed method. Finally, we conduct several ablations studies to show the significance of modules in our network.

#### 5.1. Evaluation metric

Evaluation metric is set to measure the similarity between the predicted lanes and the ground-truth lanes. Previous evaluation metric [10], which set predefined  $y$ -position to regress the  $x$  and  $z$  coordinates of lane points, is not sophisticated. Due to the fixed anchor design, this metric essentially performs badly when the lanes are horizontal.

To tackle this problem, we propose a two-stage evaluation metric, which regards 3D lane evaluation problem as a point cloud matching problem combined with top-view

constraints. Intuitively, two 3D lane pairs are well matched when achieving little difference in the  $z$ - $x$  plane (top view) alongside a close height distribution. Matching in the top view constrains the predicted lanes in the correct forward direction, and close point clouds distance in 3D space ensures the accuracy of the predicted lanes in spatial height.

Our proposed metric first calculate the matching degree of two lanes on the  $z$ - $x$  plane. To be concrete, lane is represented as  $L^k = \{(x_i^k, y_i^k, z_i^k)\}_{i=1}^n$ . To judge whether predicted lane  $L^p$  matches ground-truth lane  $L^g$ , the first matching process is done in the  $z$ - $x$  plane, namely top-view, we use the traditional IoU method [29] to judge whether  $L^p$  matches  $L^g$ . If the IoU is bigger than the IoU threshold, further we use a unilateral Chamfer Distance ( $CD$ ) to calculate the curves matching error in the camera coordinates. The curve matching error  $CD_{p,g}$  between  $L^p$  and  $L^g$  is calculated as follows:

$$\begin{cases} CD_{p,g} = \frac{1}{m} \sum_{i=1}^m \|P_{g_i} - \hat{P}_{p_j}\|_2, \\ \hat{P}_{p_j} = \min_{P_{p_j} \in L^p} \|P_{p_j} - P_{g_i}\|_2, \end{cases} \quad (11)$$

where  $P_{p_j} = (x_{p_j}, y_{p_j}, z_{p_j})$  and  $P_{g_i} = (x_{g_i}, y_{g_i}, z_{g_i})$  are point of  $L^p$  and  $L^g$  respectively, and  $\hat{P}_{p_j}$  is the nearest point to the specific point  $P_{g_i}$ .  $m$  represents the number of points taken at an equal distance from the ground-truth lane. If the unilateral chamfer distance is less than the chamfer distance threshold, written as  $\tau_{CD}$ . We consider  $L^p$  matches  $L^g$  and accept  $L^p$  as a true positive. The legend to calculate chamfer distance error is shown in Figure 5.

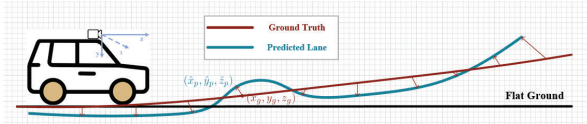


Figure 5. **Unilateral chamfer distance** Given a point on the ground-truth lane, find the nearest point on the predicted lane to calculate the chamfer distance.

This evaluation metric is intuitive and strict, and more importantly, it applies to more lane topologies such as vertical lanes, so it is more generalized. At last, since we know how to judge a predicted lane is true positive or not, we use the precision, recall and F-score as the evaluation metric.

## 5.2. Implementation details

Our experiments are carried out on our proposed *ONCE-3DLanes* benchmark. We use the Segformer [38] as our backbone with two branches. The Segformer encoder is pre-trained with Imagenet [19]. The input resolution is set to  $320 \times 800$  with our augmentation strategy during training. The augmentation is turned off during testing. The Adamw optimizer is adopted to train 20 epochs, with an

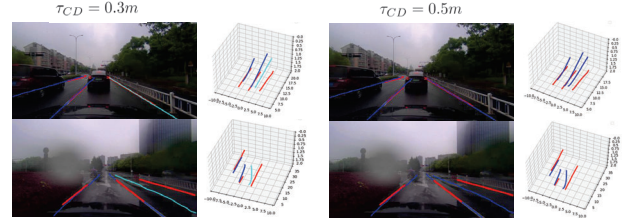


Figure 6. Qualitative results and failure cases analysis under specific threshold. The ground-truth lanes are colored in red, while true positives of our predictions are in blue and false positives in cyan. The  $\tau_{CD}$  of 0.5 is somehow too loose for discrimination.

initial learning rate at  $3e-3$  and applied with a poly scheduler by default. For evaluation, We set the IoU threshold as 0.3, the chamfer distance thresh  $\tau_{CD}$  as 0.3m. We also test our model using the Mindspore [27].

## 5.3. Benchmark performance

### 5.3.1 Main results

We train our model with the whole training set of 200k images and report the detection performance on the test set of *ONCE-3DLanes* dataset. To verify the rationality of the hyper-parameters setting in our evaluation metric, we evaluate our model under different Chamfer Distance threshold  $\tau_{CD}$  and report the testing results in Table 2.

$\tau_{CD}(m)$	F1(%)	Precision(%)	Recall(%)	CD error(m)
0.15	48.35	58.52	41.19	0.062
0.30	64.07	75.90	55.42	0.098
0.50	68.92	81.65	59.62	0.118

Table 2. Performance of *SALAD* under different  $\tau_{CD}$  thresholds on our test set.

We report our performance in different settings of  $\tau_{CD}$ , in order to fully investigate the impact of tightening or loosening the criteria on model performance. The illustrations of criteria adopting various  $\tau_{CD}$  are also proposed in Figure 6. It can be seen that under the threshold of 0.5, some predicted lanes relatively far from the ground-truth are judged as true positives. While under the  $\tau_{CD}$  of 0.15, the criteria seems too harsh. The  $\tau_{CD}$  of 0.3 is more reasonable and our *SALAD* achieves 64.07% F1-score based on this threshold. Besides, since the distance is calculated based on real scene so it is highly adaptable to real-world datasets. In the remaining parts, the experimental results are reported under the threshold of 0.3.

### 5.3.2 Results of 3D lane detection methods

In order to further verify the authenticity of our dataset and the superiority of our method, which is extrinsic-free, we

Method	<i>extri. free</i>	F1(%)	Precision(%)	Recall(%)	CD error(m)
3D LaneNet [7]	✗	44.73	61.46	35.16	0.127
GenLaneNet [10]	✗	45.59	63.95	35.42	0.121
<i>SALAD</i>	✓	<b>64.07</b>	75.90	55.42	0.098

Table 3. Performance of 3D lane detection on ONCE-3DLanes.

also conduct some experiments of other 3D lane detection algorithms on our dataset.

It is worth noting that all existing 3D lane detection algorithms need to provide camera poses as supervision information and have strict assumptions that the camera is installed at zero degrees roll relative to the ground plane [7]. However, our method requires no external parameter information. To make comparison, we used the camera-pose parameters provided by ONCE [25] to provide supervision signals for counterpart methods and ultimately evaluated them on our test set. The performance of 3D lane detection algorithms are presented in Table 3.

The comparison shows our method outperforms other 3D lane detection method in *ONCE-3DLanes* dataset. The comparison result indicates that extrinsic-required methods under the assumption of fixed camera pose and zero-degree camera roll may suffer in real 3D scenarios.

### 5.3.3 Results of extended 2D lane detection methods

*ONCE-3DLanes* dataset is a newly released 3D lane detection dataset and no previous work addresses the problem of extrinsic-free 3D lane detection. In order to verify the validity of our dataset and the efficiency of our method, we extend the existing 2D lane detection model and evaluate their performance on our dataset.

Different from 3D lane detection methods, 2D lane detection algorithms can only detect the pixel coordinates of lanes in the image plane, but cannot recover the spatial information of lanes. To obtain 3D lane detection results, we use a pre-trained depth estimation model MonoDepth2 [9] (finetuned to the depth scale of our dataset) to estimate the pixel-level depth of the image. It is worth mentioning that the depth model is finetuned on full ONCE dataset in order to avoid under-fitting caused by sparse supervision provided by lane points, which also indicates that this pipeline is difficult to be extended on other 3D lane benchmarks. Combined with the detection results of extended 2D lane detection model, the spatial position of 3D lanes are reconstructed and our evaluation metric is used for performance evaluation. The results are showed in Table 4.

Experimental results show that the extended 2D models are effective to perform 3D lane detection task on our ONCE-3DLanes dataset. It can also be found that our proposed method can reach 64.07% F1-score, outperforms the best of other methods 56.57% by 7.5%, which shows the superiority of our method.

Method	F1(%)	Precision(%)	Recall(%)	CD error(m)
PointLaneNet (R101) [3]	54.99	64.50	47.93	0.115
UltraFast (R101) [30]	54.18	63.68	47.14	0.128
RESA (R101) [40]	55.53	65.08	48.43	0.112
LaneAF (DLA34) [1]	56.39	66.07	49.18	0.109
LaneATT (R122) [35]	56.57	66.75	49.07	0.101
<i>SALAD</i>	<b>64.07</b>	75.90	55.42	0.098

Table 4. Performance of extended 2D lane detection methods on *ONCE-3DLanes* test set.

## 5.4. Ablation study

To verify the efficiency of revised data augmentation methods, we conduct ablation experiments by gradually turning off data augmentation strategy. As shown in Table 5, the performance of our method is constantly improved with the gradual introduction of our data augmentation strategy. The flip method brings a 1.06% improvement to our model and the 3D scale provides a further 1.83% improvement, which proves the validity of our augmentation strategy.

Method	F1(%)	Precision(%)	Recall(%)	CD error(m)
<i>SALAD</i> (w/o aug)	61.18	72.43	52.95	0.103
+ flip	62.24 <sup>+1.06</sup>	73.67 <sup>+1.24</sup>	53.88 <sup>+0.93</sup>	0.102 <sup>-0.001</sup>
+ joint scale	64.07 <sup>+2.89</sup>	75.90 <sup>+3.47</sup>	55.42 <sup>+2.47</sup>	0.098 <sup>-0.005</sup>

Table 5. Ablation study on data augmentation strategy.

## 6. Conclusion and limitations

In this paper, we have presented a largest real-world 3D lane detection benchmark *ONCE-3DLanes*. To revive the interest of 3D lane detection, we benchmark the dataset with a novel evaluation metric and propose an extrinsic-free and anchor-free method, dubbed *SALAD*, directly predicting the 3D lane from a single image in an end-to-end manner. We believe our work can lead to the expected and unexpected innovations in communities of both academia and industry.

As our dataset construction requires LiDAR to provide 3D information, occlusions would lead to short interruption and we have used interpolation to fix it. Missing points problem in the distance still exists due to the low resolution of LiDAR. Future work will focus on the ground point clouds completion to generate full information for 3D lanes.

**Acknowledgments** This work was supported in part by National Natural Science Foundation of China (Grant No. 6210020439), Lingang Laboratory (Grant No. LG-QS-202202-07), Natural Science Foundation of Shanghai (Grant No. 22ZR1407500), Shanghai Municipal Science and Technology Major Project (Grant No. 2018SHZDZX01 and 2021SHZDZX0103), Science and Technology Innovation 2030 - Brain Science and Brain-Inspired Intelligence Project (Grant No. 2021ZD0200204), MindSpore and CAAI-Huawei MindSpore Open Fund.



## References

- [1] Hala Abualsaud, Sean Liu, David Lu, Kenny Situ, Akshay Rangesh, and Mohan M Trivedi. Laneaf: Robust multi-lane detection with affinity fields. *arXiv preprint*, 2021. 2, 8
- [2] Luca Caltagirone, Mauro Bellone, Lennart Svensson, and Mattias Wahde. Lidar-camera fusion for road detection using fully convolutional neural networks. *Robotics and Autonomous Systems*, 2019. 2
- [3] Zhenpeng Chen, Qianfei Liu, and Chenfan Lian. Pointlanenet: Efficient end-to-end cnns for accurate real-time lane detection. In *IV*, 2019. 8
- [4] Travis J. Crayton and Benjamin Mason Meier. Autonomous vehicles: Developing a public health research agenda to frame the future of transportation policy. *Journal of Transport Health*, 2017. 1
- [5] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *SIGGRAPH*, 1996. 5
- [6] Netalee Efrat, Max Bluvstein, Noa Garnett, Dan Levi, Shaul Oron, and Bat El Shlomo. Semi-local 3d lane detection and uncertainty estimation. *arXiv preprint*, 2020. 1, 2
- [7] N. Garnett, R Cohen, T. Pe’Er, R Lahav, and L. Dan. 3d-lanenet: End-to-end 3d multiple lane detection. In *ICCV*, 2019. 1, 2, 3, 4, 5, 8
- [8] Noa Garnett, Roy Uziel, Netalee Efrat, and Dan Levi. Synthetic-to-real domain adaptation for lane detection. In *ACCV*, 2020. 1
- [9] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *ICCV*, 2019. 8
- [10] Y. Guo, G. Chen, P. Zhao, W. Zhang, J. Miao, J. Wang, and T. E. Choe. *Gen-LaneNet: A Generalized and Scalable Approach for 3D Lane Detection*. 2020. 1, 2, 3, 4, 5, 6, 8
- [11] Alberto Hata and Denis Wolf. Road marking detection using lidar reflective intensity data and its application to vehicle localization. 2014. 2
- [12] Bei He, Rui Ai, Yang Yan, and Xianpeng Lang. Accurate and robust lane detection based on dual-view convolutional neural network. In *IV*, 2016. 2
- [13] Namdar Homayounfar, Wei-Chiu Ma, Shrinidhi Kowshika Lakshmikanth, and Raquel Urtasun. Hierarchical recurrent attention networks for structured online maps, 2020. 1
- [14] Yuenan Hou, Zheng Ma, Chunxiao Liu, and Chen Change Loy. Learning lightweight lane detection cnns by self attention distillation. In *ICCV*, 2019. 2
- [15] Yujie Jin, Xiangxuan Ren, Fengxiang Chen, and Weidong Zhang. Robust monocular 3d lane detection with dual attention. In *ICIP*, 2021. 1, 2
- [16] Jiyong Jung and Sung-Ho Bae. Real-time road lane detection in urban areas using lidar data. *Electronics*, 2018. 2
- [17] Soren Kammel and Benjamin Pitzer. Lidar-based lane marker detection and mapping. In *IVS*, 2008. 2
- [18] Yeongmin Ko, Younkwan Lee, Shoaib Azam, Farzeen Munir, Moongu Jeon, and Witold Pedrycz. Key points estimation and point instance segmentation approach for lane detection. *IEEE Transactions on Intelligent Transportation Systems*, 2021. 2
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 7
- [20] Jason Ku, Ali Harakeh, and Steven L Waslander. In defense of classical image processing: Fast depth completion on the cpu. In *CRV*, 2018. 5
- [21] Xiang Li, Jun Li, Xiaolin Hu, and Jian Yang. Line-cnn: End-to-end traffic line detection with line proposal unit. *IEEE Transactions on Intelligent Transportation Systems*, 2019. 2
- [22] Lizhe Liu, Xiaohao Chen, Siyu Zhu, and Ping Tan. Condlanenet: a top-to-down lane detection framework based on conditional convolution. *arXiv preprint*, 2021. 2
- [23] Zechen Liu, Zizhang Wu, and Roland Tóth. Smoke: Single-stage monocular 3d object detection via keypoint estimation. In *CVPR workshops*, 2020. 4, 6
- [24] Sheng Lu, Zhaojie Luo, Feng Gao, Mingjie Liu, KyungHi Chang, and Changhao Piao. A fast and robust lane detection method based on semantic segmentation and optical flow estimation. *Sensors*, 2021. 2
- [25] J. Mao, M. Niu, C. Jiang, H. Liang, X. Liang, Y. Li, C. Ye, W. Zhang, Z. Li, and J. Yu. One million scenes for autonomous driving: Once dataset. 2021. 3, 8
- [26] B. Min, G. Mattyus, N. Homayounfar, S. Wang, and R. Urtasun. Deep multi-sensor lane detection. In *IROS*, 2018. 2
- [27] Mindspore. <https://www.mindspore.cn/>, 2020. 7
- [28] D. Neven, B De Brabandere, S. Georgoulis, M. Proesmans, and L Van Gool. Towards end-to-end lane detection: an instance segmentation approach. In *IV*, 2018. 1, 2, 5
- [29] Xingang Pan, Jianping Shi, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Spatial as deep: Spatial cnn for traffic scene understanding. In *AAAI*, 2018. 1, 2, 7
- [30] Zequn Qin, Huanyu Wang, and Xi Li. Ultra fast structure-aware deep lane detection. In *ECCV*, 2020. 8
- [31] Zhan Qu, Huan Jin, Yang Zhou, Zhen Yang, and Wei Zhang. Focus on local: Detecting lane marker from bottom up via key point. In *CVPR*, 2021. 2
- [32] N. E. Sela, M. Bluvstein, S. Oron, L. Dan, and B. E. Shlomo. 3d-lanenet+: Anchor free lane detection using a semi-local representation. 2020. 1, 2, 4, 5
- [33] Jinming Su, Chao Chen, Ke Zhang, Junfeng Luo, Xiaoming Wei, and Xiaolin Wei. Structure guided lane detection. *arXiv preprint arXiv:2105.05403*, 2021. 1, 2
- [34] Lucas Tabelini, Rodrigo Berriel, Thiago M Paixao, Claudine Badue, Alberto F De Souza, and Thiago Oliveira-Santos. Polylanenet: Lane estimation via deep polynomial regression. In *ICPR*, 2021. 1, 2
- [35] Lucas Tabelini, Rodrigo Berriel, Thiago M Paixão, Claudine Badue, Alberto F De Souza, and Thiago Olivera-Santos. Keep your eyes on the lane: Attention-guided lane detection. *arXiv e-prints*, 2020. 1, 2, 8
- [36] Michael Thuy and F Puente León. Lane detection and tracking based on lidar data. *Metrology and Measurement Systems*, 2010. 2

- [37] Wouter Van Gansbeke, Bert De Brabandere, Davy Neven, Marc Proesmans, and Luc Van Gool. End-to-end lane detection through differentiable least-squares fitting. In *ICCV*, 2019. [2](#)
- [38] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *arXiv preprint arXiv:2105.15203*, 2021. [4](#), [7](#)
- [39] Xinyu Zhang, Zhiwei Li, Xin Gao, Dafeng Jin, and Jun Li. Channel attention in lidar-camera fusion for lane line segmentation. *Pattern Recognition*, 2021. [2](#)
- [40] Tu Zheng, Hao Fang, Yi Zhang, Wenjian Tang, Zheng Yang, Haifeng Liu, and Deng Cai. Resa: Recurrent feature-shift aggregator for lane detection. In *AAAI*, 2021. [8](#)