Pseudo-LiDAR Based Road Detection

Libo Sun, Haokui Zhang, and Wei Yin

Abstract—Road detection is a critically important task for self-driving cars. By employing LiDAR data, recent works have significantly improved the accuracy of road detection. However, relying on LiDAR sensors limits the application of those methods when only cameras are available. In this paper, we propose a novel road detection approach with RGB images being the only input. Specifically, we exploit pseudo-LiDAR using depth estimation and propose a feature fusion network in which RGB images and learned depth information are fused for improved road detection. To optimize the network architecture and improve the efficiency of our network, we propose a method to search for the information propagation paths. Finally, to reduce the computational cost, we design a modality distillation strategy to avoid using depth estimation networks during inference. The resulting model eliminates the reliance on LiDAR sensors and achieves state-of-the-art performance on two challenging benchmarks, KITTI and R2D.

Index Terms—road detection, drivable area detection.

I. INTRODUCTION

ROAD detection is indispensable in the field of autonomous driving and advanced driver assistance systems (ADAS). As a critical component for autonomous driving, road detection aims to detect the drivable areas for autonomous vehicles to make reliable decisions. Due to the great importance of road detection, a large number of related works [1]–[6] have been proposed in recent years. More recently, with the help of deep convolutional neural networks (DCNNs) [7], [8], the performance of road detection has been significantly improved.

Introducing DCNNs can obtain accurate road detection results, but it fails when facing some extreme situations, such as ambiguous boundaries and shadow areas. To solve these problems and to improve road detection accuracy, a large number of LiDAR-based works [5], [6], [9] have been proposed. These works have demonstrated that introducing Li-DAR information can significantly increase the road detection accuracy. However, it leads to another problem that LiDAR data needs to be obtained by using expensive devices, which increases the cost on one hand and limits their application when only cameras are available on the other hand.

In order to address the above limitations, we propose a new road detection method that does not rely on LiDAR data and can achieve performance comparable to current LiDAR-based methods. Firstly, motivated by the success of using pseudo-LiDAR in other autonomous driving related tasks [13]–[15], we propose a pseudo-LiDAR based road detection framework,

L. Sun and W. Yin are with The University of Adelaide, Australia (e-mail: libo.sun@adelaide.edu.au; wei.yin@adelaide.edu.au).

H. Zhang is with School of Computer Science, Northwestern Polytechnical University, Xi'an, and also with Intellifusion, Shenzhen, China (e-mail: hkzhang1991@mail.nwpu.edu.cn).

Part of this work was carried out when H. Zhang was visiting The University of Adelaide. L. Sun and H. Zhang contributed equally to this work.


Fig. 1. Road detection results by using our method on the KITTI dataset [10]. The top and bottom parts of the figure show the results from a perspective view and an aerial view respectively. Red represents false negatives, blue areas correspond to false positives, and green denotes true positives.

in which we use pseudo-LiDAR (PL) learned from RGB images to replace LiDAR.

We further introduce a pseudo-LiDAR based information fusion network (PLIFNet) to synthesize both RGB and pseudo-LiDAR data and mutually update the two input sources. This is different from previous LiDAR-based networks [5], [6] that only use the LiDAR branch to update the RGB branch in a single direction. Secondly, the current road detection networks [5], [6] are manually designed, which are not conducive to the use of all information. To address this issue, we propose an information propagation path search (IPPS) method to automatically tune and optimize the PLIFNet architecture. The optimized PLIFNet can obtain state-of-the-art performance on both the KITTI and R2D datasets without using LiDAR data (see Fig. 1 for an example). Lastly, to reduce the computational cost of inference, we propose a modality distillation method for the road detection task. Particularly, a new modality distillation loss is introduced to learn expressive features of different granularities from pixel inputs to patches and holistic images. By doing so, it enables our framework to obtain accurate prediction results without using computationally costly depth networks during inference.

Our main contributions can be summarized as follows:

- We propose a pseudo-LiDAR based information fusion framework, in which we use pseudo-LiDAR learned from RGB images to replace LiDAR and introduce a new information fusion network, called PLIFNet, to fuse RGB and pseudo-LiDAR information.
- To further enhance the information fusion network, we propose an information propagation path search approach
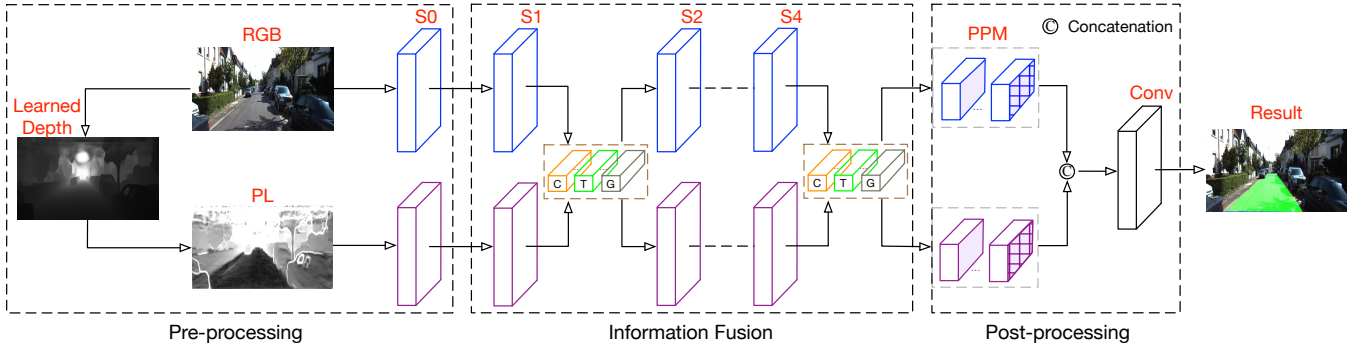
Fig. 2. The overall pipeline of the proposed method. The whole framework consists of three sub-components, i.e., pre-processing, information fusion, and post-processing. The residual blocks in ResNet [11] are denoted as S0 to S4, and the RGB and depth features are fused in a CTG module (see Fig. 3) from S1 to S4. The RGB and pseudo-LiDAR (PL) information are combined to obtain the final prediction result after the pyramid pooling module (PPM) [12]. Different from the information fusion approach used in LiDAR-based networks [5], [6], two branches exchange their information in the CTG modules and the searched information propagation paths (see Fig. 4) are added to compose the final network.

(IPPS) to automatically tune and optimize our manually designed PLIFNet. With RGB images being the only input, our optimized network outperforms the current state-of-the-art RGB and LiDAR-based methods.

- We propose a modality distillation strategy to free our PLIFNet from additional computational costs due to the generation of pseudo-LiDAR, making the pipeline of using PLIFNet more concise and efficient.

II. RELATED WORK

Our model is related to road detection, neural architecture search (NAS), and knowledge distillation. Thus, we focus on discussing related works in these three areas.

A. Road Detection

Traditional road detection methods [3], [4], [16], [17] are usually based on edges, colors, and textures. Alvarez and Lopez [16] introduced illuminant-invariant features which can be combined with a road class-likelihood classifier to obtain the road detection results. He et al. [3] presented an approach in which boundaries are estimated on intensity images and road areas are detected using the full color images. Alvarez et al. [4] proposed a method to estimate road priors and combine different contextual cues, including horizon lines, vanishing points, lane markings, 3D scene layout, and road geometry, to obtain detection results. Jeong et al. [18] proposed a classification method for road detection in which images are divided into sub-images, and a K-means-based classifier is used to perform classification for sub-images. Kumar et al. [19] designed a method to incorporate contextual cues to find lanes between lane curves and improve the computational efficiency without compromising accuracy. These traditional methods are proposed to solve similar problems as ours, i.e., detecting drivable regions for autonomous vehicles. However, traditional methods are generally less effective than current learning-based methods in terms of accuracy, as shown in the recent literature [5], [6], [20], and thus, we mainly focus on learning-based methods.

With the development of convolutional neural networks (CNNs), learning-based approaches are widely used in the field of semantic segmentation [21]–[27]. Based on these semantic segmentation networks, a large number of CNN-based road detection methods [1], [5], [6], [20], [28] have been proposed for the road detection task. There is a strong correlation between semantic segmentation and road detection, as both of them can be regarded as pixel-wise labeling tasks. However, compared with semantic segmentation, road detection only focuses on road regions and has higher requirements for prediction performance.

Qi et al. [29] introduced a learning-based method to fuse RGB images, contours, and location priors in a network to predict road detection results. Chen and Chen [1] proposed a network in which they formulate the road detection and road boundary detection problem into a unified Bayesian network model to detect roads and road boundaries. Sun et al. [20] proposed a residual learning-based network with a residual refinement module composed of reverse and boundary attention units to perform road detection. Chen et al. [5] presented an approach to transform LiDAR data space to RGB data space to combine two types of information together to obtain road detection results. Compared with the methods that only use RGB images, recent LiDAR-based methods [5], [6] have shown obvious accuracy advantages, which inspired us to propose our pseudo-LiDAR based framework.

B. Neural Architecture Search

Neural architecture search (NAS) aims to discover high-performance neural architectures, enabling researchers to discard the tedious and heuristic manual neural architecture design process. A widespread phenomenon is that architectures built by NAS algorithms show better performance than manually designed architectures. Evolutionary algorithms (EAs) are used by early methods for optimizing neural architectures and parameters. The best architecture can be obtained by iteratively mutating a set of candidate architectures [30]. An alternative to EAs is using reinforcement learning (RL) methods, e.g., policy gradients [31] and Q-learning [32], to train a recurrent

neural network as a meta-controller to generate potential architectures. Recently, a large number of speed-up methods [33]–[36] have been proposed to accelerate search speed, and NAS algorithms have been introduced into various tasks. For image classification, based on the continuous relaxation of the architecture representation, Liu et al. [36] proposed DARTS, in which gradient information is used to obtain the most useful path. For semantic segmentation, Auto-DeepLab [37] expanded the search space of DARTS by introducing multiple paths with different widths and resolutions. For image restoration, Suganuma et al. proposed E-CAE [38], where EA is adopted to search the architecture of convolutional auto-encoders for image inpainting and denoising.

In this work, we propose an NAS algorithm to search for information propagation paths to optimize the road detection network. The proposed IPPS is mostly related to DARTS [36]. Both of them are gradient-based NAS algorithms, but the search space and target are different. DARTS is proposed to search for a classification network. The IPPS is proposed to further optimize the road detection network by finding valuable paths. The search space of DARTS is on all candidate operations, while the search space of our IPPS is on all candidate information propagation paths. To our knowledge, our IPPS is the first NAS work that is proposed to fine-tune the information propagation paths in a manually designed network.

C. Knowledge Distillation

In the work of [39], knowledge distillation is introduced by distilling the knowledge in an ensemble of models into a single model to obtain results on MNIST. In recent years, a large number of knowledge distillation methods [40]–[48] have been proposed. Knowledge distillation are widely used in different types of tasks, such as object detection [40], [49], image classification [50], [51], and semantic segmentation [52]. Chen et al. [40] proposed an end-to-end approach to learn compact multi-class object detection models through knowledge distillation. Park et al. [50] presented a relational knowledge distillation method and showed its application in image classification and metric learning. Tang et al. [53] proposed ranking distillation, which is the first knowledge distillation approach used in recommender systems. Pan et al. [54] designed an enhanced collaborative denoising auto-encoder (ECAE) model for recommender systems, in which a knowledge distillation is used to learn useful knowledge and reduce noise. Garcia et al. [55] proposed a modality distillation approach to teach a hallucination network to mimic the depth stream in a multiple stream action recognition network. These distillation approaches can make networks to maintain their effectiveness on platforms with limited computing or input resources, which has significantly expanded their application scope.

Inspired by these distillation works, we design a modality distillation approach to reduce the number of parameters and computational costs of our framework. We notice that a similar modality distillation idea has been used in CMT-CNN [56]. However, CMT-CNN is proposed for pedestrian detection while our method is proposed for road detection. Different
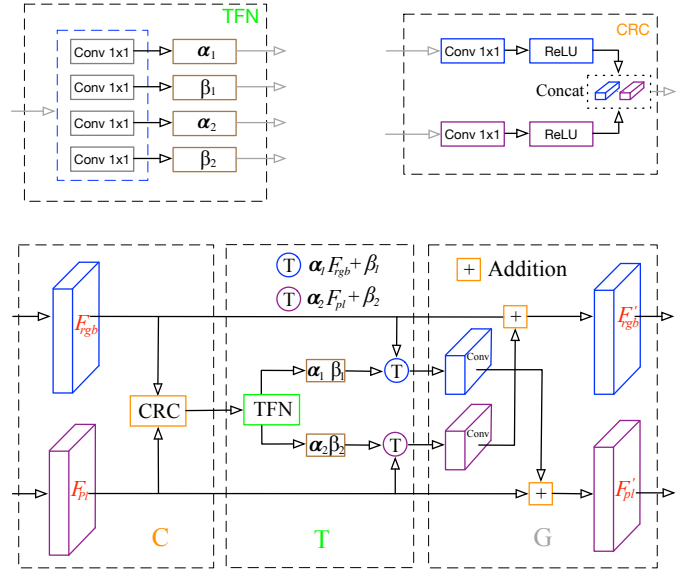

Fig. 3. The detailed information fusion module CTG. The CTG module consists of three parts, i.e., concatenation (C), transformation (T), and generation (G). Firstly, RGB and PL features are concatenated together after the $1{\times}1$ convolution and ReLU operations (top right). Then, the concatenated features are fed into a 'TFN' module (top left) to obtain the α and β parameters for fusion. Finally, the updated features are generated based on α and β.

network architectures are also independently used in the two methods to meet different task requirements. During inference, CMT-CNN does not need to input thermal maps, but it relies on a region reconstruction network (RRN), while our method does not have such a requirement.

III. OUR METHOD

In this section, we elaborate on our pseudo-LiDAR based road detection framework, which is inspired by previous pseudo-LiDAR based works [13]–[15] and LiDAR-based works [5], [6]. To begin with, we introduce our framework to fuse RGB and pseudo-LiDAR information in Section III-A. Then we explain how to automatically tune and optimize the manually designed road detection network using the proposed IPPS in Section III-B. Finally, in Section III-C we present our modality distillation strategy, which is used to free our network from additional computation costs due to the use of depth estimation networks.

A. Pseudo-LiDAR Based Road Detection Framework

The proposed pseudo-LiDAR based road detection framework is shown in Fig. 2. It consists of three sub-components: pre-processing, information fusion, and post-processing. First, pseudo-LiDAR is generated in the pre-processing step. Then, the RGB and PL information are fused in the information fusion component. Finally, the RGB and PL features are combined together in the post-processing component to output the final detection results.

Pre-processsing. Current state-of-the-art methods [5], [6] usually need to use LiDAR data to improve their performance. However, as LiDAR data must be collected using specific

devices, these methods are no longer applicable when only cameras are available. To eliminate the requirement on the use of LiDAR data, we introduce PL in our framework. For each RGB image, we first use a depth estimation network to generate its depth. Then we transform the learned depth into PL. In our framework, the BTS [57] network is used for depth estimation. Before training the road detection network, the depth network is trained independently using ground truth depth. When training the road detection network, the depth network is only used to generate PL and will not be trained again. Note that the BTS network can be replaced by any other monocular depth estimation networks [58]–[63].

Similar to the LiDAR-based work [5], we transform all the depth values to their attitude space to generate PL. For each pixel (u, v) in a depth image, its attitude value $H_{(u,v)}$ in the camera coordinate system can be derived as follows:

$$H_{(u,v)} = \frac{(v - c_v)}{f_v} \times z \quad \text{with} \quad z = D_{(u,v)}, \qquad (1)$$

where $D_{(u,v)}$ represents the depth value of pixel (u, v), and f_v and c_v are camera focal length and the v component of the principal point position (c_u, c_v).

After transforming all the depth values to attitude values, the $PL(u, v)$ which describes the pseudo-LiDAR value at coordinate position (u, v) is defined as:

$$PL_{(u,v)} = \frac{1}{N} \sum_{N_u, N_v} \sqrt{\left(\frac{\delta_h}{N_u - u}\right)^2 + \left(\frac{\delta_h}{N_v - v}\right)^2} \qquad (2)$$
$$\text{with} \quad \delta_h = H_{(N_u, N_v)} - H_{(u,v)},$$

where (N_u, N_v) represents neighbourhood positions of (u, v) and N is the total number of neighbourhood positions. The (N_u, N_v) is set as positions within a 7×7 window centered at (u, v).

Information Fusion. We use ResNet [11] as the backbone to build a two-branch fusion structure. The main layers in the ResNet are denoted as stage 1 to stage 4. RGB features and PL features are fused to exchange their information after each stage. The detailed process to fuse and update RGB and PL features after each stage is shown in Fig. 3. The proposed information fusion module CTG consists of three sub-components, i.e., concatenation (C), transformation (T), and generation (G). In component C, a Conv-ReLU-Concat (CRC) module is constructed to concatenate features from two branches. Due to the difference between RGB information and PL information, the features extracted from RGB images are inconsistent with the features extracted from PL. To solve this feature space inconsistency problem, the RGB and PL features are transformed to each other's data space in the information fusion process. The data space transformation is treated as a linear transformation, and the component T is used for generating transformation parameters α_1, β_1, α_2, and β_2. In component T, a transformation network (TFN) inputs the merged features and outputs the transformation parameters. Finally, the fused features are generated in component G as follows:

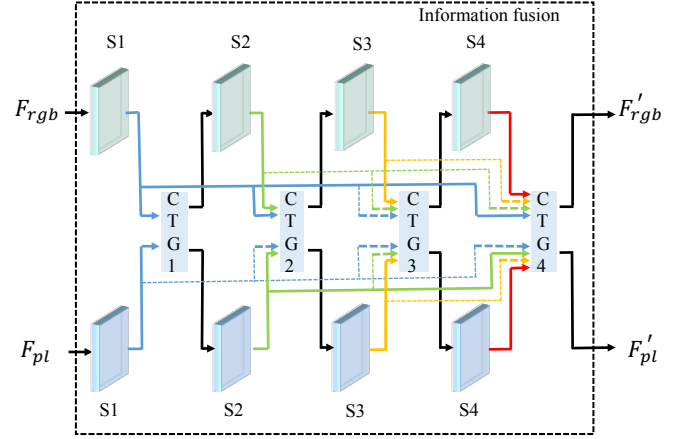$$F'_{rgb} = F_{rgb} + (\alpha_1 F_{pl} + \beta_1), \qquad (3)$$

Fig. 4. Searching for information propagation paths in the information fusion component. Arrows with colors of blue, green, yellow, and red denote features from stage 1 to 4 respectively. Arrows in bold represent the paths searched by our proposed IPPS algorithm.

$$F'_{pl} = F_{pl} + (\alpha_2 F_{rgb} + \beta_2), \qquad (4)$$

where F_{rgb} and F'_{rgb} represent the RGB features before and after information fusion respectively, while F_{pl} and F'_{pl} represent the PL features before and after information fusion respectively.

Post-processing. After four stages of information exchange and fusion, features from the RGB branch and the PL branch are combined in the post-processing component to predict the final road detection results. In the post-processing component, the RGB and PL features are sent to the pyramid pooling module (PPM) [12] to fuse features from different pyramid scales. For the details of PPM operation, we refer to [12]. Then, the RGB and PL information are combined together using a concatenation operation. Finally, classification convolutions (Conv) [12] are performed to predict results.

B. Information Propagation Path Search

In deep learning models, the features from shallow layers contain rich detailed information and the features from deep layers reveal rich semantic information. Fusing shallow features and deep features is beneficial for keeping output details. Currently, most of the dense prediction networks need to combine features from different layers to improve the prediction accuracy. As road detection needs to label each pixel in an image, it can also be regarded as a dense prediction task. Therefore, we can further improve the performance of the proposed road detection network by adding information propagation paths between shallow layers and deep layers.

Adding information propagation paths can combine information from shallow layers and deep layers, but it causes a problem as to how to find the paths. Furthermore, adding too many paths will introduce unnecessary parameters and make the network overloaded. On the other hand, adding too few paths cannot make full use of the shallow and deep features.

To solve this problem, we propose an information propagation path search (IPPS) algorithm to search for the most valuable paths. As shown in Fig. 4, we first build an overloaded

information fusion module by adding all the possible information propagation paths in the original information fusion module as presented in Section III-A. Specifically, in the original information fusion module, each CTG module only takes features from its current stage as input, while in the overloaded information fusion module, each CTG module takes not only the features from its current stage but also the features from all the previous stages as input. We denote the CTG module in stage i as CTG_i. In the original CTG module, the output of CTG_i is calculated as:

$$h_i^{rgb}, h_i^{pl} = CTG_i(s_i^{rgb}, s_i^{pl}), \tag{5}$$

where s_i^{rgb} and s_i^{pl} represent the output of the RGB channel and the PL channel of the ith stage. h_i^{rgb} and h_i^{pl} are the outputs of CTG_i. During the searching process, the output of CTG_i is:

$$h_i^{rgb}, h_i^{pl} = CTG_i(o_i^{rgb}, o_i^{pl}),$$
$$o_i^c = \sum_{j=1}^{i} p_{c,i}^{j \to i} T_{c,i}^{j \to i}(s_j^c), \tag{6}$$

where p is a set of weights of the information propagation paths. $T_{c,i}^{j \to i}$ denotes the transfer function between the jth stage and the ith stage in channel c. $T_{c,i}^{j \to i}$ consists of a single convolution layer, which is used to make the dimension and spatial resolution of s_j^c consistent with that of s_i^c. When j is equal to i, it is an identity mapping. The purpose of searching information propagation paths is to learn p, which is updated via gradient descent during the searching process. The optimization function of the searching process can be expressed as:

$$\arg\min_{p,\theta} Segloss(F(x, y|(p, \theta))), \tag{7}$$

where $Segloss$ is the cross entropy loss for road detection, $F()$ denotes the road detection network, x and y are the input image and the corresponding ground truth label respectively, and θ represents the kernels of convolution layers of the network. During the searching process, p and θ are updated alternately. After optimization, for each stage, we keep the path which has the maximum weight. If the maximum weight is negative, we will discard all paths from shallow layers to the current stage.

C. Modality Distillation

By using the architecture optimization of IPPS, our framework can achieve state-of-the-art performance. However, generating PL relies on depth estimation networks and introduces additional computation costs. To reduce computational costs and make our framework more concise during inference, we propose a modality distillation (MD) method.

As shown at the top of Fig. 5, our MD framework consists of two main parts, the trained teacher part and the student part. Two channels of the trained teacher network take RGB and PL information as inputs. In contrast, both of the two channels of the student network take RGB images as inputs. During training, the student network is trained to learn the
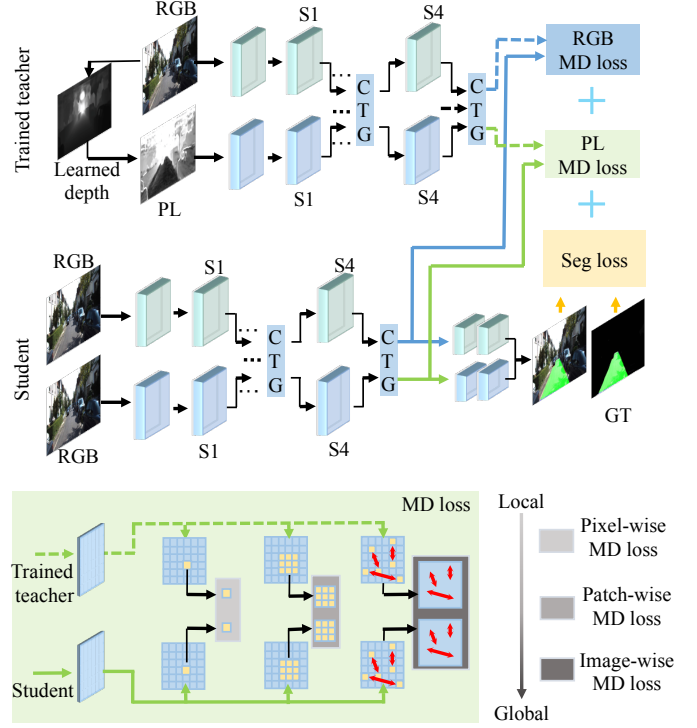

Fig. 5. Modality distillation (MD) framework and MD loss. During inference, the student network only needs to take RGB images as inputs. From local to global, our proposed MD loss consists of three parts, i.e., pixel-wise loss, patch-wise loss, and image-wise loss.

features from the teacher network by using different losses. The optimization function is defined as:

$$L_{MD} = \lambda l_{MD,RGB} + \lambda l_{MD,PL} + l_{seg},$$
$$l_{MD,RGB} = l_{pi}(h_t^1, h_s^1) + l_{pa}(h_t^1, h_s^1) + l_{im}(h_t^1, h_s^1), \tag{8}$$
$$l_{MD,PL} = l_{pi}(h_t^2, h_s^2) + l_{pa}(h_t^2, h_s^2) + l_{im}(h_t^2, h_s^2),$$

where h_t^1 and h_s^1 represent the outputs of channel 1 of the last CTG modules of the trained teacher and student networks respectively. h_t^2 and h_s^2 denote the outputs of channel 2 of the last CTG modules of the trained teacher and student networks respectively. l_{seg} is the segmentation loss, which is defined as the cross entropy loss between the prediction and ground truth. $l_{MD,RGB}$ and $l_{MD,PL}$ are the modality distillation losses of the RGB and PL channels respectively. From local to global, each MD loss consists of three items: pixel-wise MD loss (l_{pi}), patch-wise MD loss (l_{pa}), and image-wise MD loss (l_{im}).

Pixel-wise MD loss. This is the loss item that restricts the distance between features for pixel pairs. We use the squared difference to formulate this loss:

$$l_{pi}(h_t, h_s) = \frac{\sum_{i=1}^{H} \sum_{j=1}^{W} \left(f(h_t^{ij}) - f(h_s^{ij}) \right)^2}{H \times W} \tag{9}$$

where H and W denote the spatial resolution of the feature map. $f()$ is the sigmoid function. h_t^{ij} is the feature vector at location (i, j) of the trained teacher network, and h_s^{ij} is the feature vector at location (i, j) of the student network.

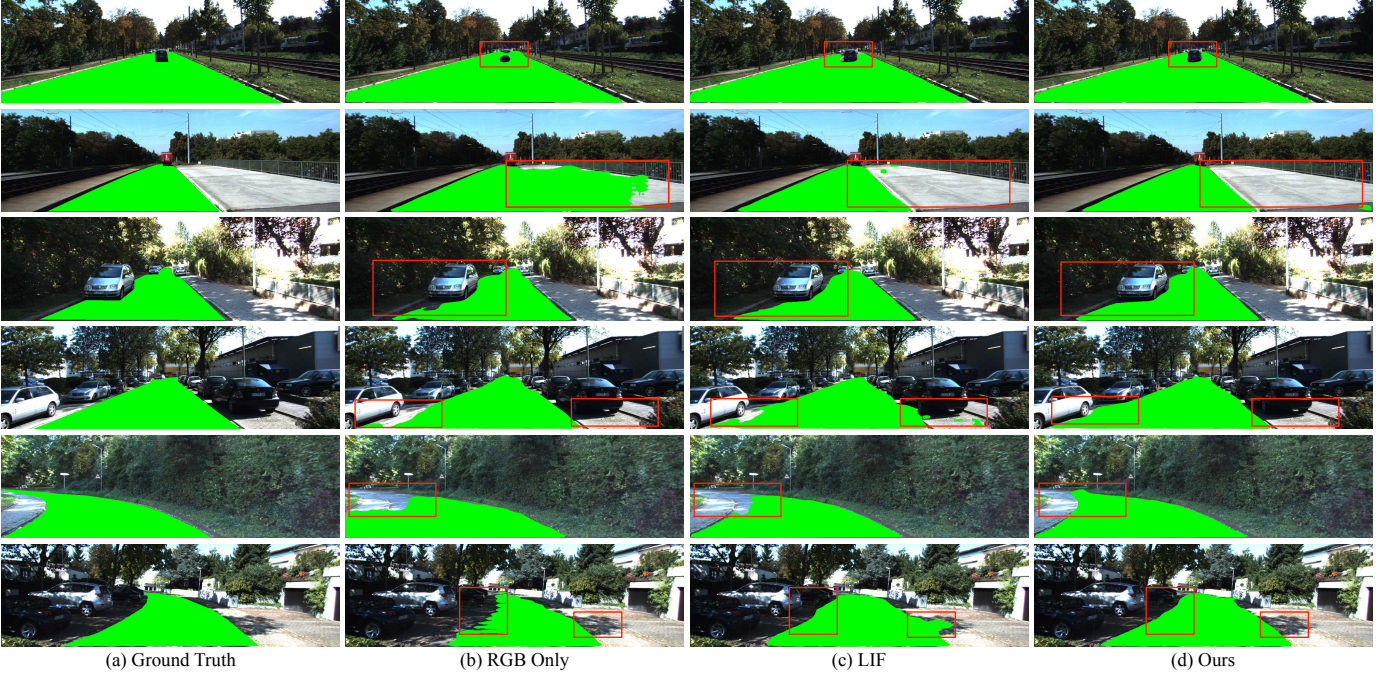(a) Ground Truth (b) RGB Only (c) LIF (d) Ours

Fig. 6. Visualization comparisons on the KITTI dataset. The comparison details are marked with red boxes. Three types of visualization results are presented, i.e., RGB only, LIF [5], and our optimized PLIF.

TABLE I
BENEFITS OF PL BASED INFORMATION FUSION. 'NF' DENOTES NO
INFORMATION FUSION; 'LIF' REPRESENTS THE INFORMATION FUSION
APPROACH USED IN [5]; 'PLIF' INDICATES OUR PROPOSED PL BASED
INFORMATION FUSION. RESULTS ARE BASED ON 5-FOLD
CROSS-VALIDATION ON THE KITTI DATASET. BEST SCORES ARE IN BOLD.

Method	Input	MaxF (%)	AP (%)	PRE (%)
NF	RGB	94.88	92.19	94.97
NF	Depth	95.07	92.19	95.44
NF	PL	95.30	91.95	95.44
LIF	RGB + PL	95.44	92.26	95.60
PLIF	RGB + PL	**96.06**	**92.30**	**95.95**

TABLE II
COMPARISON BETWEEN USING IPPS PATHS AND USING ALL
PROPAGATION PATHS.

Method	MaxF (%)	AP (%)	PRE (%)	REC (%)
All Path	95.25	92.09	95.24	95.26
IPPS Path	**96.23**	**92.22**	**96.14**	**96.32**

TABLE III
PERFORMANCE COMPARISON BETWEEN USING GROUND TRUTH (GT)
DEPTH AND LEARNED DEPTH (LD). RESULTS ARE BASED ON THE R2D
DATASET AND BEST SCORES ARE IN BOLD.

Input	ACC (%)	F-Score (%)	PRE (%)	REC (%)
LD	98.72	98.42	97.70	99.15
GT	**99.30**	**99.13**	**98.90**	**99.36**

TABLE IV
EFFECTIVENESS OF MODALITY DISTILLATION. THREE COMPONENTS ARE
EVALUATED, I.E., PI (PIXEL-WISE MD LOSS), PA (PATCH-WISE MD
LOSS), AND IM (IMAGE-WISE MD LOSS). RESULTS ARE BASED ON
5-FOLD CROSS-VALIDATION ON THE KITTI DATASET. BEST SCORES ARE
IN BOLD.

PI	IM	PA	MaxF (%)	AP (%)	PRE (%)	REC (%)
-	-	-	94.72	92.09	94.75	94.70
✓	-	-	95.10	92.19	95.19	95.01
✓	✓	-	95.23	92.17	95.21	95.25
✓	✓	✓	**95.41**	**92.32**	**95.55**	**95.26**

Patch-wise MD loss. This is the loss item that restricts the structure similarity of feature patches. This loss is calculated as:

$$l_{pa}(h_t, h_s) = \log_{10}(\text{ssim}(h_t, h_s)^{-1}), \quad (10)$$

where ssim() is structural similarity [64]. The window size of ssim() is set to 11 in this paper.

Image-wise MD loss. This is the loss term that is used to maintain structure consistency of the whole feature map. The loss function is defined as:

$$l_{im}(h_t, h_s) = \frac{1}{n} \sum_{r=1}^{n} \frac{\cos^{-1}(\text{Sim}(h'_{t,r}, h'_{s,r}))}{\pi}, \quad (11)$$

where n is the number of random sampling, and Sim() means similarity, which is defined as:

$$\text{Sim}(h'_{t,r}, h'_{s,r}) = \frac{h'_{t,r} \cdot h'_{s,r}}{||h'_{t,r}|| \, ||h'_{s,r}||},$$
$$h'_{t,r} = h_t^{i_r j_r} - h_t^{i'_r j'_r},$$
$$h'_{s,r} = h_s^{i_r j_r} - h_s^{i'_r j'_r}, \quad (12)$$

where $\cdot$ is inner product and $|| \; ||$ represents the norm of a vector. (i_r, j_r) and (i'_r, j'_r) are two sets of coordinate positions which are randomly selected in the rth sampling.

TABLE V
USING DIFFERENT PSEUDO-LIDAR REPRESENTATIONS. 'PLXYZ'
REPRESENTS USING THE PSEUDO-LIDAR IN 3D OBJECTION
METHODS [13], [14]. 'PLDY' REPRESENTS THE PSEUDO-LIDAR IN OUR
FRAMEWORK.

Transformation	MaxF (%)	AP (%)	PRE (%)	REC (%)
PLXYZ	95.62	**92.25**	95.60	95.65
PLDY	**96.23**	92.22	**96.14**	**96.32**

TABLE VI
BENEFITS OF USING MODALITY DISTILLATION. 'DEPTHNET' AND
'ROADNET' DENOTE THE NETWORK TO GENERATE DEPTH AND OUR ROAD
DETECTION NETWORK RESPECTIVELY. 'FULL' REPRESENTS OUR FULL
METHOD AND 'MD' DENOTES THE PROPOSED MODALITY DISTILLATION
METHOD.

Metric	DepthNet	RoadNet	Full	MD
Speed	0.21 s	0.25 s	0.21 + 0.25 s	0.25 s
Params	112.8 M	97.89 M	112.8 + 97.89 M	97.89 M

IV. EXPERIMENTS

We perform comprehensive ablation studies and comparisons on the KITTI [10] and R2D [6] datasets to investigate the effectiveness of our method. First, to investigate the effectiveness of each part in our method, we adapt 5-fold cross-validation to perform ablation studies on the KITTI dataset. Then, we compare our method with other state-of-the-art methods on both the KITTI and R2D datasets.

A. Datasets

The KITTI road benchmark is one of the most popular road detection datasets. The KITTI road detection dataset contains three different categories of road scenes, i.e., urban unmarked roads (UM), urban multiple marked roads (UMM), and urban unmarked roads (UU). To avoid overfitting and carry out fair comparisons, the number of submissions is strictly restricted on the KITTI server. Thus, we investigate the effectiveness of each part of the proposed method by using 5-fold cross-validation on the KITTI training dataset, and only the final full method is evaluated on the KITTI benchmark server.

The R2D dataset is a large-scale synthetic road detection dataset. It is collected in various weather conditions and contains 11,430 images with corresponding depth images and semantic labels. The authors of the R2D dataset split the whole dataset into three subsets, i.e., 6117 training images, 2624 validation images, and 2689 testing images. As the R2D dataset contains multiple scenarios in different environments, it can be used to evaluate the effectiveness of methods in different road scenes.

B. Evaluation Metrics

For the evaluation on the KITTI dataset, we use the standard evaluation metrics provided by the KITTI benchmark. The metrics consist of maximum F1-measure (MaxF), precision rate (PRE), average precision (AP), false positive rate (FPR), false negative rate (FNR), and recall rate (REC). For more details about these six metrics, we refer to [10].
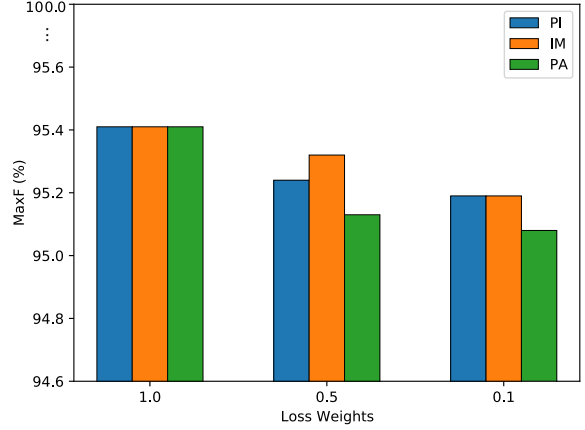

Fig. 7. The performance of using different loss weights on the KITTI dataset. PI, PA, and IM represent pixel-wise loss, patch-wise loss, and image-wise loss respectively. To reduce the number of variables, each time only one loss weight variable is changed, while the other two are fixed to 1.0. For example, if the weight of PA is changed to 0.1, the weights of PI and IM are set to 1.0.

Fig. 8. Road detection results on the KITTI benchmark [10]. Top: urban marked (UM), middle: urban unmarked (UU), bottom: urban multiple marked (UMM).

Similar to the evaluation on the KITTI dataset, five metrics are used by R2D to evaluate the performance. The metrics include accuracy (ACC), precision (PRE), recall (REC), F-score, and intersection over union (IoU). The details of the evaluation metrics on R2D are described in [6]. For all the experiments on the R2D dataset, we follow the same evaluation approach as used in [6].

C. Implementation Details

We implement our network using PyTorch [69] and train it on NVIDIA V100 GPUs. To obtain the high-quality depth
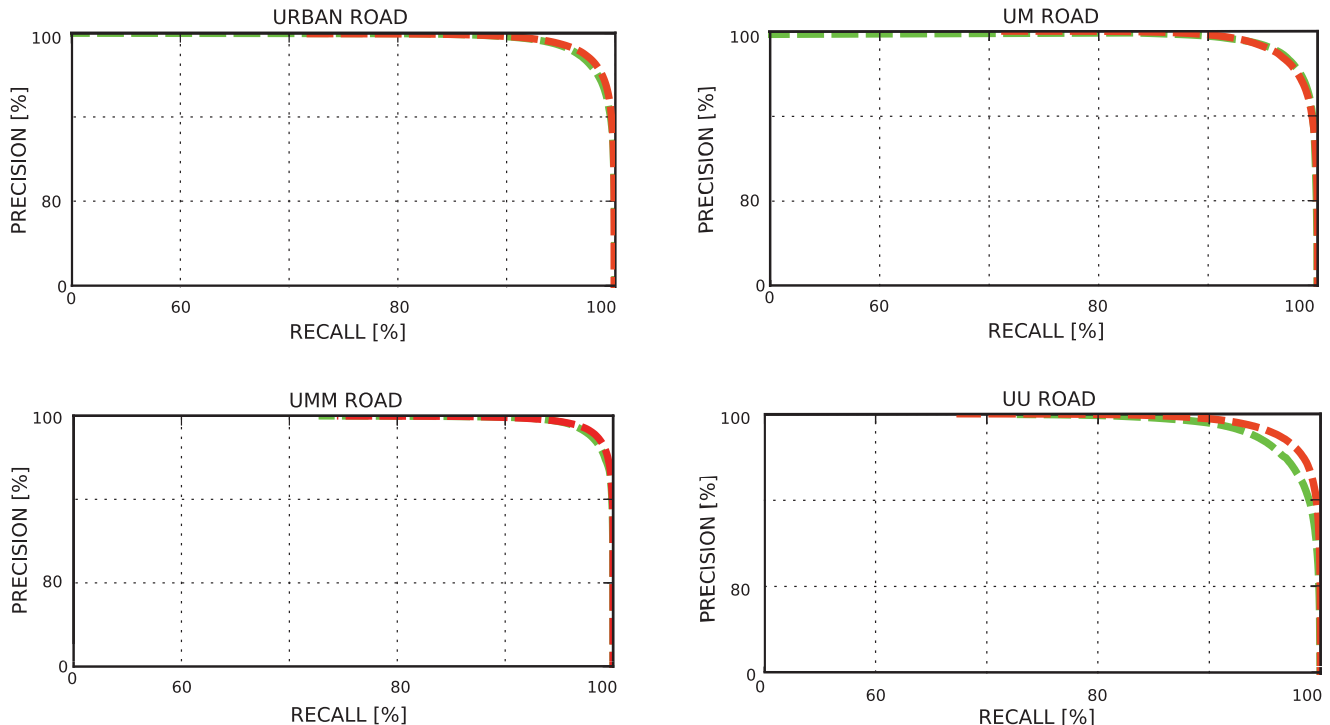
Fig. 9. Precision and recall curves on the KITTI benchmark. Three categories are presented, i.e., urban marked roads (UM), urban multiple marked roads (UMM) and urban unmarked roads (UU). Urban road is a combination of these three categories. Two methods are compared, i.e., PLARD (green lines) and ours (red lines).

TABLE VII
COMPARISON WITH THE TOP-10 PUBLISHED ROAD DETECTION METHODS ON THE KITTI BENCHMARK (URBAN_ROAD RANKING). BEST SCORES ARE IN BOLD. 'FULL' DENOTES USING THE PROPOSED FULL INFORMATION FUSION NETWORK IN WHICH THE LEARNED DEPTH IS USED TO GENERATE PL.

Methods	Input	MaxF (%)	AP (%)	PRE (%)	REC (%)	FPR (%)	FNR (%)	Rank
RBNet [1]	RGB	94.97	91.49	94.94	95.01	2.79	4.99	11
TVFNet [65]	RGB + LiDAR	95.34	90.26	95.73	94.94	2.33	5.06	10
SSLGAN [2]	RGB	95.53	90.35	95.84	95.24	2.28	4.76	9
RGB36-Cotrain [66]	RGB	95.55	93.71	95.68	95.42	2.37	4.58	8
LC-CRF [67]	RGB + LiDAR	95.68	88.34	93.62	**97.83**	3.67	**2.17**	7
NIM-RTFNet [68]	RGB + LiDAR	96.02	94.01	96.43	95.62	1.95	4.38	6
LidCamNet [9]	RGB + LiDAR	96.03	93.93	96.23	95.83	2.07	4.17	5
RBANet [20]	RGB	96.30	89.72	95.14	97.50	2.75	2.50	4
SNE-RoadSeg [6]	RGB + LiDAR	96.75	94.07	96.90	96.61	1.70	3.39	3
PLARD [5]	RGB + LiDAR	97.03	94.03	97.19	96.88	1.54	3.12	2
Ours (Full)	RGB	**97.42**	**94.09**	**97.30**	97.54	**1.49**	2.46	1

to generate PL, the BTS [57] network is used for monocular depth estimation. The BTS network can be replaced by any other monocular depth estimation networks for PL generation. The experimental settings are similar in the ablation studies and the final evaluations. During training and testing, the input image size is resized to 384×1280 for the KITTI dataset and 480×640 for the R2D dataset. We use SGD to optimize all the parameters and the batch size is set to 4 on both the KITTI and R2D datasets. To better evaluate the effectiveness of each component of our network, models are trained from scratch without data augmentation in our ablation studies. For the final evaluations on the KITTI and R2D datasets, parameters from the model of [5] and our KITTI model are used as the initial weights for the training respectively. The initial learning rates are set to 1×10^{-2} in our ablation studies and modality

distillation. In addition, for the final benchmark evaluation, we use several data augmentation approaches to train our models, including random cropping and multi-scale training and testing.

D. Ablation Studies

The Effectiveness of Information Fusion. In this section, different approaches are compared to investigate the effectiveness of our information fusion approach, including NF, LIF, and our proposed PLIF. NF represents using only RGB, depth, or PL information as input. LIF means using both RGB and PL information, and only the information from the PL branch is fused into the RGB branch. PLIF denotes our proposed information fusion module, in which not only the PL information is fused into the RGB branch but also the RGB in-
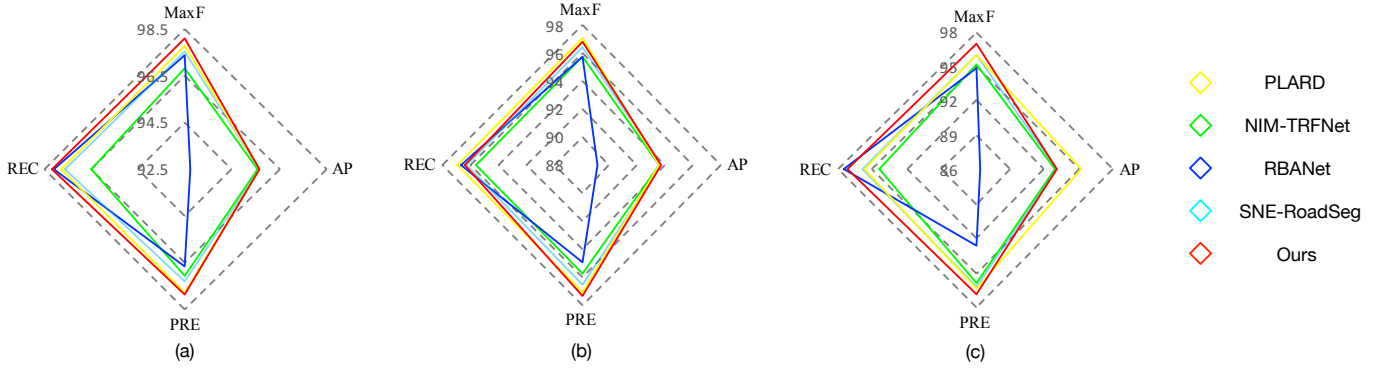
Fig. 10. The performance of different categories on the testing set of the KITTI benchmark. (a) urban multiple marked roads (UMM); (b) urban marked roads (UM); (c) urban unmarked roads (UU).

TABLE VIII
COMPARISONS OF UU, UM, AND UMM ON THE KITTI BENCHMARK. BEST 'MAXF' SCORES ARE IN BOLD. 'FULL' DENOTES USING THE PROPOSED
FULL INFORMATION FUSION NETWORK IN WHICH THE LEARNED DEPTH IS USED TO GENERATE PL.

Method	Input	UU		UM		UMM	
		MaxF (%)	AP (%)	MaxF (%)	AP (%)	MaxF (%)	AP (%)
RBNet [1]	RGB	93.21	89.18	94.77	91.42	96.06	93.49
TVFNet [65]	RGB + LiDAR	93.65	87.57	94.96	89.17	96.47	93.16
SSLGAN [2]	RGB	94.40	87.84	94.62	89.50	96.72	92.99
RGB36-Cotrain [66]	RGB	94.53	92.54	94.55	93.12	96.75	95.39
LC-CRF [67]	RGB + LiDAR	94.01	85.24	94.91	86.41	97.08	92.06
NIM-RTFNet [68]	RGB + LiDAR	95.11	92.94	95.71	93.56	96.79	95.61
LidCamNet [9]	RGB + LiDAR	94.54	92.74	95.62	93.54	97.08	95.51
RBANet [20]	RGB	94.91	86.35	95.78	89.14	97.38	92.67
SNE-RoadSeg [6]	RGB + LiDAR	96.03	93.03	96.42	93.67	97.47	95.63
PLARD [5]	RGB + LiDAR	95.95	95.25	**97.05**	93.53	97.77	95.64
Ours (Full)	RGB	**96.93**	93.08	96.87	93.71	**98.05**	95.63

formation is fused into the PL branch. Results of 5-fold cross-validation on the KITTI dataset are listed in Table I. From the comparison results in Table I, we can observe: (1) Using PL information can obtain a comparable performance with using RGB information. Thus, the PL information should be considered equally with the RGB information. Compared with using depth information as input, taking the PL information derived from depth information achieves better performance; (2) By considering PL and RGB information equally and making better use of all the information, our proposed PLIF achieves better performance than the LIF proposed in [5]. Using our PLIF approach increases the MaxF by 0.62% compared with using the LIF approach.

Benefits of Using IPPS. We conduct experiments on the KITTI dataset to evaluate the effectiveness of the proposed IPPS method. The comparison results are shown in Table II. For the MaxF metric, we can observe that using all the information propagation paths only gets 95.25% while using the IPPS path obtains 96.23%. The main reason for this is that not all the paths are valuable for information fusion. The proposed IPPS method enables the network to select useful information propagation paths from all the propagation paths to obtain better accuracy.

Visualization Comparisons. In addition to the comparisons using performance metrics, visualization comparisons are also used to investigate the effectiveness of our method. For a better visualization comparison, some visualization results are high-lighted in Fig. 6. As can be seen from Fig. 6, compared with using RGB information only or the LiDAR-based information fusion (LIF) network [5], our proposed PLIF can obtain better prediction results. For example, in some boundary and shadow areas, our prediction results are closer to the ground truth. This is mainly because our PLIF network can make more effective use of RGB and pseudo-LiDAR information.

Effects of Using Learned Depth. As we use a depth estimation network to obtain learned depth to build our pseudo-LiDAR based framework, we perform an analysis to investigate the difference between using ground truth depth and learned depth. The analysis is performed on the R2D dataset, as the R2D datastet is a synthetic dataset and can provide enough reliable dense depth. We can observe from Table III that using ground truth depth to replace learned depth can obtain better performance, but the performance cannot be significantly improved. There are two reasons for this situation. The first reason is that it is difficult to obtain further significant improvements, as our network has achieved very high performance. Another reason is that RGB information and pseudo-LiDAR information work together in our framework, which can make the network robust to slight depth errors.

The Effectiveness of Modality Distillation. We perform comprehensive experiments to illustrate the effectiveness of our modality distillation approach on the KITTI dataset. The comparison results are shown in Table IV. As can be seen from Table IV, from local to global, with the modality distillation

loss items being added one by one, the accuracy of modality distillation shows steady growth. We can observe that by using all three proposed MD loss items, our modality distillation strategy increases MaxF from 94.72% to 95.41%.

Effects of Using Different Loss Weights. Because different losses are used to supervise the training of our modality distillation model, we conduct experiments to investigate the impact of using different loss weights. To reduce the complexity for our analysis, each time we only change one loss weight and fix the other two. The results of using different loss weights are shown in Fig. 7. From Fig. 7, we can observe that reducing any loss weight can result in a decrease in accuracy. Therefore, three loss weights are set to equal factors of 1.0 in the modal distillation.

Benefits of Modality Distillation. To reduce parameters and computational costs, we propose a modality distillation approach. The comparison between using modality distillation and without using modality distillation is shown in Table VI. From the comparison, we observe that the modality distillation strategy can significantly reduce the number of parameters and improve processing speed. Specifically, to generate PL, we use BTS-ResNext101 for depth estimation. The BTS-ResNext101 has 112.8 M parameters, and each inference takes 0.21 seconds on NVIDIA GTX 1060. However, for our road detection network, the number of parameters and the time used for each inference are about 100 M and 0.25 seconds. Thus, the proposed modality distillation method can reduce nearly 50% parameters and running time. Introducing modality distillation achieves a good trade-off between performance and computational costs.

Using Different Pseudo-LiDAR Representations. Our framework is based on the pseudo-LiDAR generation as detailed in Section III-A. Because transforming depth into 3D LiDAR points has already been used in the field of 3D object detection [13]–[15], we conduct experiments to investigate the performance of using pseudo-LiDAR transformation adopted by 3D object detection methods. We follow the implementation of [13] to transform depth into the LiDAR points for a comparison. As shown in Table V, introducing pseudo-LiDAR representation used by 3D object detection methods shows worse performance and decreases the MaxF from 96.23% to 95.61%. Therefore, the pseudo-LiDAR representation for 3D object detection is not used in our road detection framework.

E. Comparison with State-of-the-art Methods

Evaluation on the KITTI Benchmark. We test our method on the KITTI benchmark to evaluate its performance. Some qualitative results on the testing set of KITTI benchmark are shown in Fig. 8. Precision and recall (PR) curves [70], [71] which show the relationship between precision and recall are obtained on the KITTI dataset [10] to evaluate the accuracy. The PR curve comparisons between our method and PLARD [5] are shown in Fig. 9. From these curves, we can observe that for some categories, such as UU, our method can even obtain better precision if the same recall rate is used. The reason for this is that our framework is more effective for information fusion. The proposed PLIF approach in which the

pseudo-LiDAR information and RGB information are used to update each other is more effective than the information fusion approach used in PLARD. We also note that our method lightly falls behind PLARD on UM. This is because the UM, UMM, and UU are combined together as a whole dataset (Urban) for training. Training the network using all the categories usually sacrifices performance in some subcategories. There is also a similar phenomenon in semantic segmentation, that is, a method obtains better overall performance, but shows slightly worse performance in some subcategories.

We compare our models with other top ten methods published on the KITTI benchmark. The urban road detection results which can be regarded as the overall performance of all the categories are shown in Table VII. It can be seen that our method outperforms all the other approaches on urban road detection by using only RGB images as input. Except for the overall performance, we also compare the results on three categories: UMM, UM, and UU in Table VIII. From Table VIII, we notice that our method obtains the best results on UU and UMM and the second best result on UM. Our method improves the best MaxF from 96.03% to 96.63% (+0.9%) on UU. To make a more intuitive comparison, a visualized metric comparison is shown in Fig. 10. From the comparison results of these three categories, we can observe that our method shows better overall performance than other methods, i.e., it outperforms other methods on most of the metrics. Unlike previous methods which only use LiDAR information as an auxiliary to update the RGB branch, the proposed PLIF architecture enables the RGB branch and pseudo-LiDAR branch to exchange their information at different stages. Therefore, our method is more effective for information fusion to obtain better performance. In addition, we also compare our modality distillation results with other image-based top performing methods on the KITTI benchmark. The comparison is shown in Table IX. Benefiting from the proposed modality distillation strategy, our network can achieve highly comparable performance without using depth networks during inference.

Evaluation on the R2D Dataset. We evaluate our method on the R2D dataset to demonstrate its effectiveness with different illumination and weather conditions. The visualization results of our method on the R2D dataset are shown in Fig. 11. As can be seen from Fig. 11, our method is robust and can obtain accurate results in different weather conditions, such as rain days and sunny days. It demonstrates that the proposed framework is effective and robust in diverse scenarios. Similar to the comparison on the KITTI dataset, we also use accuracy metrics to compare our method with other methods on the R2D dataset. As the R2D dataset is a new dataset that has just been proposed in [6], we only compare the results which are available on the R2D dataset with ours. Unlike the KITTI dataset, the R2D dataset does not provide a sever to submit and evaluate results. Therefore, for a fair comparison, we use the evaluation code of [6] to compute all the metrics. The detailed comparison results on the R2D dataset are shown in Table X. As can be seen from Table X, our method outperforms other state-of-the-art approaches on the R2D dataset by using only RGB images. Compared with 'Ours (Full)' which takes RGB

TABLE IX
THE MODALITY DISTILLATION RESULTS AND THE COMPARISON WITH OTHER METHODS ON THE KITTI BENCHMARK. ONLY THE METHODS WITHOUT USING ANY REFINEMENT OPERATION ARE COMPARED. 'MD' REPRESENTS USING THE PROPOSED MODALITY DISTILLATION APPROACH (WITHOUT USING THE LEARNED DEPTH TO GENERATE PL IN TESTING).

Methods	MaxF (%)	AP (%)	PRE (%)	REC (%)	FPR (%)	FNR (%)
RBNet [1]	94.97	91.49	94.94	95.01	2.79	4.99
SSLGAN [2]	95.53	90.35	95.84	95.24	2.28	4.76
RGB36-Cotrain [66]	95.55	93.71	95.68	95.42	2.37	4.58
Ours (MD)	**95.79**	**93.95**	**96.12**	**95.46**	**2.12**	**4.54**

TABLE X
PERFORMANCE COMPARISON ON THE R2D DATASET. THE RESULTS OF SNE-ROADSEG ARE FROM [6]. 'FULL' REPRESENTS OUR FULL METHOD (LEANED DEPTH IS USED TO GENERATE THE PL). 'MD' REPRESENTS USING THE PROPOSED MODALITY DISTILLATION APPROACH (WITHOUT USING THE LEARNED DEPTH TO GENERATE THE PL DURING INFERENCE).

| Methods	Backbone	Input	ACC (%)	PRE (%)	REC (%)	F-Score (%)	IoU (%)
SNE-RoadSeg [6]	ResNet-101	RGBD	98.0	98.2	97.1	97.6	95.4
SNE-RoadSeg [6]	ResNet-152	RGBD	98.6	99.1	97.6	98.3	96.7
Ours (MD)	ResNet-101	RGB	99.1	99.2	98.5	98.9	97.8
Ours (Full)	ResNet-101	RGB	**99.5**	**99.5**	**99.3**	**99.4**	**98.8**

Fig. 11. Prediction results on the R2D dataset. 'Full' represents using our full framework. 'MD' denotes using the proposed modality distillation method. All the images are resized for better visualization.

and learned PL as inputs, the modality distillation strategy slightly reduces the performance. However, the performance using modality distillation is still better than that of SNE-RoadSeg [6].

V. CONCLUSION

In this paper, we propose a novel pseudo-LiDAR based road detection framework. The framework is proposed to obtain state-of-the-art performance without using LiDAR. To achieve this goal, we use pseudo-LiDAR learned from RGB images to replace LiDAR and design a new pseudo-LiDAR based information fusion network to fuse RGB and pseudo-LiDAR information. To further optimize the proposed framework, we propose an IPPS algorithm to automatically search for the valuable information propagation paths to optimize our manually designed network. Compared with previous RGB and LiDAR-based methods, the optimized network can obtain state-of-the-art performance on the KITTI and R2D datasets without using LiDAR information. Finally, we propose a modality distillation approach which makes our method to eliminate the additional parameters and computational costs due to the generation of pseudo-LiDAR from RGB images during inference.

REFERENCES

[1] Z. Chen and Z. Chen, "RBNet: A deep neural network for unified road and road boundary detection," in *Adv. Neural Inform. Process. Syst.*, 2017, pp. 677–687.

[2] X. Han, J. Lu, C. Zhao, S. You, and H. Li, "Semisupervised and weakly supervised road detection based on generative adversarial networks," *IEEE Sign. Process. Letters*, vol. 25, no. 4, pp. 551–555, 2018.

[3] Y. He, H. Wang, and B. Zhang, "Color-based road detection in urban traffic scenes," *IEEE Trans. Intell. Transp. Syst.*, vol. 5, no. 4, pp. 309–318, 2004.

[4] J. M. Alvarez, A. M. López, T. Gevers, and F. Lumbreras, "Combining priors, appearance, and context for road detection," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 3, pp. 1168–1178, 2014.

[5] Z. Chen, J. Zhang, and D. Tao, "Progressive LiDAR adaptation for road detection," *IEEE/CAA J. Autom. Sinica.*, vol. 6, no. 3, pp. 693–702, 2019.

[6] R. Fan, H. Wang, P. Cai, and M. Liu, "SNE-RoadSeg: Incorporating surface normal information into semantic segmentation for accurate freespace detection," in *Eur. Conf. Comput. Vis.*, 2020, pp. 340–356.

[7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Adv. Neural Inform. Process. Syst.*, vol. 25, pp. 1097–1105, 2012.

[8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[9] L. Caltagirone, M. Bellone, L. Svensson, and M. Wahde, "LIDAR–camera fusion for road detection using fully convolutional neural networks," *Robotics & Autonomous Syst.*, vol. 111, pp. 125–131, 2019.

[10] J. Fritsch, T. Kuehnl, and A. Geiger, "A new performance measure and evaluation benchmark for road detection algorithms," in *IEEE Conf. Intell. Transp. Syst.*, 2013, pp. 1693–1700.

[11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778.

[12] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 2881–2890.

[13] Y. Wang, W.-L. Chao, D. Garg, B. Hariharan, M. Campbell, and K. Q. Weinberger, "Pseudo-LiDAR from visual depth estimation: Bridging the gap in 3D object detection for autonomous driving," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 8445–8453.

[14] X. Weng and K. Kitani, "Monocular 3D object detection with pseudo-LiDAR point cloud," in *Int. Conf. Comput. Vis. Worksh.*, 2019, pp. 0–0.

[15] Y. You, Y. Wang, W.-L. Chao, D. Garg, G. Pleiss, B. Hariharan, M. Campbell, and K. Q. Weinberger, "Pseudo-LiDAR++: Accurate depth for 3D object detection in autonomous driving," *arXiv: Comp. Res. Repository*, vol. abs/1906.06310, 2019.

[16] J. M. Á. Alvarez and A. M. Ĺopez, "Road detection based on illuminant invariance," *IEEE Trans. Intell. Transp. Syst.*, vol. 12, no. 1, pp. 184–193, 2010.

[17] P. Lombardi, M. Zanin, and S. Messelodi, "Switching models for vision-based on-board road detection," in *IEEE Conf. Intell. Transp. Syst.*, 2005, pp. 67–72.

[18] P. Jeong and S. Nedevschi, "Efficient and robust classification method using combined feature vector for lane detection," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 15, no. 4, pp. 528–537, 2005.

[19] R. K. Satzoda and M. M. Trivedi, "On enhancing lane estimation using contextual cues," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 25, no. 11, pp. 1870–1881, 2015.

[20] J.-Y. Sun, S.-W. Kim, S.-W. Lee, Y.-W. Kim, and S.-J. Ko, "Reverse and boundary attention network for road segmentation," in *Int. Conf. Comput. Vis. Worksh.*, 2019, pp. 0–0.

[21] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 3431–3440.

[22] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, 2017.

[23] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Eur. Conf. Comput. Vis.*, 2018, pp. 801–818.

[24] H. Zhao, X. Qi, X. Shen, J. Shi, and J. Jia, "ICNet for real-time semantic segmentation on high-resolution images," in *Eur. Conf. Comput. Vis.*, 2018, pp. 405–420.

[25] F. Meng, K. Luo, H. Li, Q. Wu, and X. Xu, "Weakly supervised semantic segmentation by a class-level multiple group cosegmentation and foreground fusion strategy," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 30, no. 12, pp. 4823–4836, 2019.

[26] J. Ji, R. Shi, S. Li, P. Chen, and Q. Miao, "Encoder-decoder with cascaded CRFs for semantic segmentation," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 31, no. 5, pp. 1926–1938, 2020.

[27] F. Lv, G. Lin, P. Liu, G. Yang, S. J. Pan, and L. Duan, "Weakly-supervised cross-domain road scene segmentation via multi-level curriculum adaptation," *IEEE Trans. Circuit Syst. Video Technol.*, 2020.

[28] M. Bucher, T.-H. Vu, M. Cord, and P. Pérez, "Zero-shot semantic segmentation," *Adv. Neural Inform. Process. Syst.*, vol. 32, pp. 468–479, 2019.

[29] Q. Wang, J. Gao, and Y. Yuan, "Embedding structured contour and location prior in siamesed fully convolutional networks for road detection," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 1, pp. 230–241, 2017.

[30] H. Liu, K. Simonyan, O. Vinyals, C. Fernando, and K. Kavukcuoglu, "Hierarchical representations for efficient architecture search," *arXiv: Comp. Res. Repository*, vol. abs/1711.00436, 2017.

[31] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 8697–8710.

[32] Z. Zhong, J. Yan, W. Wu, J. Shao, and C.-L. Liu, "Practical block-wise neural network architecture generation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 2423–2432.

[33] C. Zhang, M. Ren, and R. Urtasun, "Graph hypernetworks for neural architecture search," *arXiv: Comp. Res. Repository*, vol. abs/1810.05749, 2018.

[34] T. Elsken, J. H. Metzen, and F. Hutter, "Efficient multi-objective neural architecture search via Lamarckian evolution," *arXiv: Comp. Res. Repository*, vol. abs/1804.09081, 2018.

[35] H. Pham, M. Y. Guan, B. Zoph, Q. V. Le, and J. Dean, "Efficient neural architecture search via parameter sharing," *arXiv: Comp. Res. Repository*, vol. abs/1802.03268, 2018.

[36] H. Liu, K. Simonyan, and Y. Yang, "DARTS: Differentiable architecture search," *arXiv: Comp. Res. Repository*, vol. abs/1806.09055, 2018.

[37] C. Liu, L.-C. Chen, F. Schroff, H. Adam, W. Hua, A. L. Yuille, and L. Fei-Fei, "Auto-DeepLab: Hierarchical neural architecture search for semantic image segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 82–92.

[38] M. Suganuma, M. Ozay, and T. Okatani, "Exploiting the potential of standard convolutional autoencoders for image restoration by evolutionary search," in *Int. Conf. Mach. Learn.*, 2018.

[39] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv: Comp. Res. Repository*, vol. abs/1503.02531, 2015.

[40] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker, "Learning efficient object detection models with knowledge distillation," in *Adv. Neural Inform. Process. Syst.*, 2017, pp. 742–751.

[41] J. Yim, D. Joo, J. Bae, and J. Kim, "A gift from knowledge distillation: Fast optimization, network minimization and transfer learning," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 4133–4141.

[42] R. Yu, A. Li, V. I. Morariu, and L. S. Davis, "Visual relationship detection with internal and external linguistic knowledge distillation," in *Int. Conf. Comput. Vis.*, 2017, pp. 1974–1982.

[43] S. H. Lee, D. H. Kim, and B. C. Song, "Self-supervised knowledge distillation using singular value decomposition," in *Eur. Conf. Comput. Vis.*, 2018, pp. 335–350.

[44] X. Wang, R. Zhang, Y. Sun, and J. Qi, "KDGAN: Knowledge distillation with generative adversarial networks." in *NIPS*, 2018, pp. 783–794.

[45] J. H. Cho and B. Hariharan, "On the efficacy of knowledge distillation," in *Int. Conf. Comput. Vis.*, 2019, pp. 4794–4802.

[46] F. Tung and G. Mori, "Similarity-preserving knowledge distillation," in *Int. Conf. Comput. Vis.*, 2019, pp. 1365–1374.

[47] X. Jin, B. Peng, Y. Wu, Y. Liu, J. Liu, D. Liang, J. Yan, and X. Hu, "Knowledge distillation via route constrained optimization," in *Int. Conf. Comput. Vis.*, 2019, pp. 1345–1354.

[48] B. Peng, X. Jin, J. Liu, D. Li, Y. Wu, Y. Liu, S. Zhou, and Z. Zhang, "Correlation congruence for knowledge distillation," in *Int. Conf. Comput. Vis.*, 2019, pp. 5007–5016.

[49] Y. Hao, Y. Fu, Y.-G. Jiang, and Q. Tian, "An end-to-end architecture for class-incremental object detection with knowledge distillation," in *Int. Conf. Multimedia and Expo*, 2019, pp. 1–6.

[50] W. Park, D. Kim, Y. Lu, and M. Cho, "Relational knowledge distillation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 3967–3976.

[51] Y. Liu, L. Sheng, J. Shao, J. Yan, S. Xiang, and C. Pan, "Multi-label image classification via knowledge distillation from weakly-supervised detection," in *ACM Int. Conf. Multimedia*, 2018, pp. 700–708.

[52] Y. Liu, C. Shu, J. Wang, and C. Shen, "Structured knowledge distillation for dense prediction," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020.

[53] J. Tang and K. Wang, "Ranking distillation: Learning compact ranking models with high performance for recommender system," in *Int. Conf. Knowledge Discovery & Data Mining*, 2018, pp. 2289–2298.

[54] Y. Pan, F. He, and H. Yu, "A novel enhanced collaborative autoencoder with knowledge distillation for top-N recommender systems," *Neurocomputing*, vol. 332, pp. 137–148, 2019.

[55] N. C. Garcia, P. Morerio, and V. Murino, "Modality distillation with multiple stream networks for action recognition," in *Eur. Conf. Comput. Vis.*, 2018, pp. 103–118.

[56] D. Xu, W. Ouyang, E. Ricci, X. Wang, and N. Sebe, "Learning cross-modal deep representations for robust pedestrian detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 5363–5371.

[57] J. H. Lee, M.-K. Han, D. W. Ko, and I. H. Suh, "From big to small: Multi-scale local planar guidance for monocular depth estimation," *arXiv: Comp. Res. Repository*, vol. abs/1907.10326, 2019.

[58] Y. Cao, Z. Wu, and C. Shen, "Estimating depth from monocular images as classification using deep fully convolutional residual networks," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 28, no. 11, pp. 3174–3182, 2017.

[59] H. Yan, X. Yu, Y. Zhang, S. Zhang, X. Zhao, and L. Zhang, "Single image depth estimation with normal guided scale invariant deep convolutional fields," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 29, no. 1, pp. 80–92, 2017.

[60] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 270–279.

[61] H. Mohaghegh, N. Karimi, S. R. Soroushmehr, S. Samavi, and K. Najarian, "Aggregation of rich depth-aware features in a modified stacked generalization model for single image depth estimation," *IEEE Trans. Circuit Syst. Video Technol.*, vol. 29, no. 3, pp. 683–697, 2018.

[62] K. Zheng, Z.-J. Zha, Y. Cao, X. Chen, and F. Wu, "LA-Net: Layout-aware dense network for monocular depth estimation," in *ACM Int. Conf. Multimedia*, 2018, pp. 1381–1388.

[63] C. Wang, X. Chen, S. Min, and F. Wu, "Structure generation and guidance network for unsupervised monocular depth estimation," in *Int. Conf. Multimedia and Expo*. IEEE, 2019, pp. 1264–1269.

[64] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004.

[65] S. Gu, Y. Zhang, J. Yang, J. M. Alvarez, and H. Kong, "Two-view fusion based convolutional neural network for urban road detection," in *IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2019, pp. 6144–6149.

[66] L. Caltagirone, L. Svensson, M. Wahde, and M. Sanfridson, "Lidar-camera co-training for semi-supervised road detection," *arXiv: Comp. Res. Repository*, vol. abs/1911.12597, 2019.

[67] S. Gu, Y. Zhang, J. Tang, J. Yang, and H. Kong, "Road detection through CRF based LiDAR-camera fusion," in *Int. Conf. Robotics & Automation*, 2019, pp. 3832–3838.

[68] H. Wang, R. Fan, Y. Sun, and M. Liu, "Applying surface normal information in drivable area and road anomaly detection for ground mobile robots," *arXiv: Comp. Res. Repository*, vol. abs/2008.11383, 2020.

[69] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "PyTorch: An imperative style, high-performance deep learning library," *Adv. Neural Inform. Process. Syst.*, vol. 32, pp. 8026–8037, 2019.

[70] J. Davis and M. Goadrich, "The relationship between precision-recall and ROC curves," in *Int. Conf. Mach. Learn.*, 2006, pp. 233–240.

[71] D. M. Powers, "Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation," *arXiv: Comp. Res. Repository*, vol. abs/2010.16061, 2020.

Libo Sun is currently a Ph.D. candidate with the School of Computer Science, The University of Adelaide, Australia. Before that, he worked as a research associate at Nanyang Technological University, Singapore. His current research interests include deep learning and its related applications in autonomous driving and robotics.

Haokui Zhang received the Ph.D. degree and the M.S. degree in computer application technology from Shannxi Provincial Key Laboratory of Speech and Image Information Processing, China, in 2021 and 2016, respectively. He is currently working as an algorithm researcher in Intellifusion, Shenzhen. His research interests cover information retrieval, image restoration, and hyperspectral image classification.

Wei Yin is a Ph.D. student at School of Computer Science, The University of Adelaide, Australia. He obtained his B.S. and M.Sc. from Nanjing University of Aeronautics and Astronautics. His research interests include monocular depth estimation and semantic segmentation. He also does some works in zero-shot transfer learning.