

BEVSegFormer: Bird's Eye View Semantic Segmentation From Arbitrary Camera Rigs

Lang Peng¹, Zhirong Chen¹, Zhangjie Fu¹, Pengpeng Liang², Erkang Cheng^{*1}

¹Nullmax ²Zhengzhou University

{penglang, chenzhirong, fuzhangjie, chengerkang}@nullmax.ai, liangpcs@gmail.com

Abstract

Semantic segmentation in bird's eye view (BEV) is an important task for autonomous driving. Though this task has attracted a large amount of research efforts, it is still challenging to flexibly cope with arbitrary (single or multiple) camera sensors equipped on the autonomous vehicle. In this paper, we present BEVSegFormer, an effective transformer-based method for BEV semantic segmentation from arbitrary camera rigs. Specifically, our method first encodes image features from arbitrary cameras with a shared backbone. These image features are then enhanced by a deformable transformer-based encoder. Moreover, we introduce a BEV transformer decoder module to parse BEV semantic segmentation results. An efficient multi-camera deformable attention unit is designed to carry out the BEV-to-image view transformation. Finally, the queries are reshaped according to the layout of grids in the BEV, and upsampled to produce the semantic segmentation result in a supervised manner. We evaluate the proposed algorithm on the public nuScenes dataset and a self-collected dataset. Experimental results show that our method achieves promising performance on BEV semantic segmentation from arbitrary camera rigs. We also demonstrate the effectiveness of each component via ablation study.

1. Introduction

Bird's-eye-view (BEV) representation of perception information is critical in the autonomous driving or robot navigation system as it is convenient for planning and control tasks. For example, in a mapless navigation solution, building a local BEV map provides an alternative to High-definition map (HD map) and is important for downstream tasks of the perception system including behavior prediction of agents and motion planning. BEV semantic segmentation from cameras is usually treated as the first

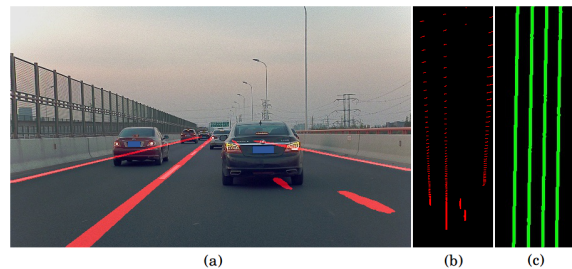


Figure 1. Comparison of lane segmentation results on image space and BEV space. (a) lane segmentation on image space, (b) BEV segmentation by IPM view transformation of (a), (c) our BEV lane segmentation.

step to build the local BEV map.

To obtain BEV semantic segmentation from cameras, traditional methods usually generate segmentation results in the image space and then convert it to the BEV space by inverse perspective mapping (IPM) function. Although IPM is a straightforward and simple way to bridge the image space and the BEV space, it requires accurate intrinsic and extrinsic parameters of the camera or real-time camera pose estimation. Therefore, it is likely to produce inferior view transformation. Taking lane segmentation as an example, as shown in Fig. 1, traditional methods with IPM provide inaccurate results in challenging scenarios where occlusions are present or at the distant areas.

Recently, deep learning approaches have been studied for BEV semantic segmentation [11, 19–23, 26]. Lift-Splat-Shoot [21] completes the view transformation from image view to BEV with pixel-wise depth estimation results. Using depth estimation increases the complexity of the view transformation process. Some approaches apply MLP [20] or FC [23] operators to perform the view transformation. These fixed view transformation methods learn a fixed mapping between the image space and the BEV space, and therefore are not dependent on the input data.

Transformer-based approaches are another line of research for perception in BEV space. In object detection task, DETR3D [29] introduces a 3D bounding boxes detec-

*Corresponding author.

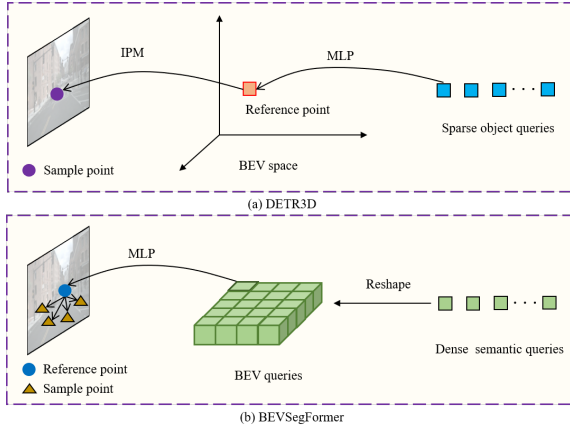


Figure 2. Comparison of view transformation between our BEV-to-image scheme and 3D-to-2D method in DETR3D [29].

tion method that directly generates predictions in 3D space from 2D features of multiple camera images. The view transformation between 3D space and 2D image space is achieved by 3D-to-2D sparse queries of a cross-attention module.

Inspired by DETR3D [29], we propose a method to compute the view transformation by BEV-to-image queries using the cross-attention mechanism in the transformer. With the novel view transformation approach, we build a BEV semantic segmentation method, BEVSegFormer, to perform BEV semantic segmentation from arbitrary camera configurations. As shown in Fig. 2, our approach differs from DETR3D [29] in several ways. First, DETR3D constructs a set of sparse object queries, while our BEVSegFormer builds dense BEV queries for the semantic segmentation task. Second, in DETR3D, a query applies MLP to predict a reference point in 3D space and then projects it back to image feature space by IPM which requires camera extrinsic parameters. In contrast, a query in our method directly predicts a reference point on image feature space through MLP operator. In this way, the view transformation of our method does not rely on camera extrinsic parameters. In addition, in order to encode more image context features for a query, our method also uses deformable attention to regress sampled points around the reference point.

Our BEVSegFormer consists of three major components: (1) a shared backbone to extract feature maps of arbitrary cameras; (2) a transformer encoder to embed feature maps by a self-attention module; and (3) a BEV transformer decoder to process BEV queries by cross-attention mechanism and output the final BEV semantic segmentation result. In the BEV transformer decoder, we introduce a multi-camera deformable cross-attention module to link feature maps from multiple cameras to BEV queries, without the need of camera intrinsic and extrinsic parameters. Specifically, a BEV query pays more attention to

the learned reference points and the corresponding sampled points in multi-camera features. In this way, the multi-camera deformable cross-attention module provides an effective method to complete the BEV-to-image view transformation.

In the experiments, we evaluate our BEVSegFormer on the public nuScenes dataset and a self-collected dataset. Our proposed BEVSegFormer sets a new state-of-the-art for BEV segmentation on the nuScenes validation set without using temporal information. The effectiveness of each proposed component is validated as well.

2. RELATED WORK

Semantic Segmentation. Semantic segmentation plays a vital role in high-level scene understanding and is a fundamental problem in computer vision. With the rapid development of CNNs, recent methods like FCN [27] and U-Net [24] apply the encoder-decoder architecture to learn the dense prediction with feature extraction backbones (e.g., VGGNet [28] and ResNet [9]). Existing methods also make use of strategies to further improve the segmentation performance, including atrous convolution [4], pyramid pooling module [34], chained residual pooling mechanism [14] and so on. In this paper, we apply ResNet as part of the encoder to extract semantic features.

BEV Semantic Segmentation. In autonomous driving and robot navigation, BEV semantic segmentation is an important perception task for downstream functions, such as behavior prediction and planning. Cam2BEV [22] performs a spatial transformer module to transform perspective features to BEV space from surrounding inputs by IPM, which is a straightforward way to link image space to BEV under flat ground assumption. SBEVNet [8] first uses stereo image features to generate disparity feature volume to enhance the BEV representation, and then estimates the BEV semantic layout on the BEV representation via a U-Net [24] model. Methods in [11, 19, 21, 26] utilize depth information to perform the view transformation. For example, Lift-Splat-Shoot [21] first estimates implicit pixel-wise depth information and then uses camera geometry to build the connection between BEV segmentation and feature maps. FIERY [11] takes a similar approach for video based input. BEV-Seg [19] uses an additional parse network to refine the projected and incomplete BEV segmentation by depth result. VPN [20] uses MLP to directly generate BEV segmentation from cross-view images. FishingNet [10] extends it to support multiple sensors. PyrOccNet [23] applies FC operators to complete the transformation at different distance ranges. Similarly, HDMapNet [12] also applies FC operators to build the local BEV map from multiple cameras. Different from previous works, Monolayout [18] performs view transformation via a standard encoder-decoder structure.

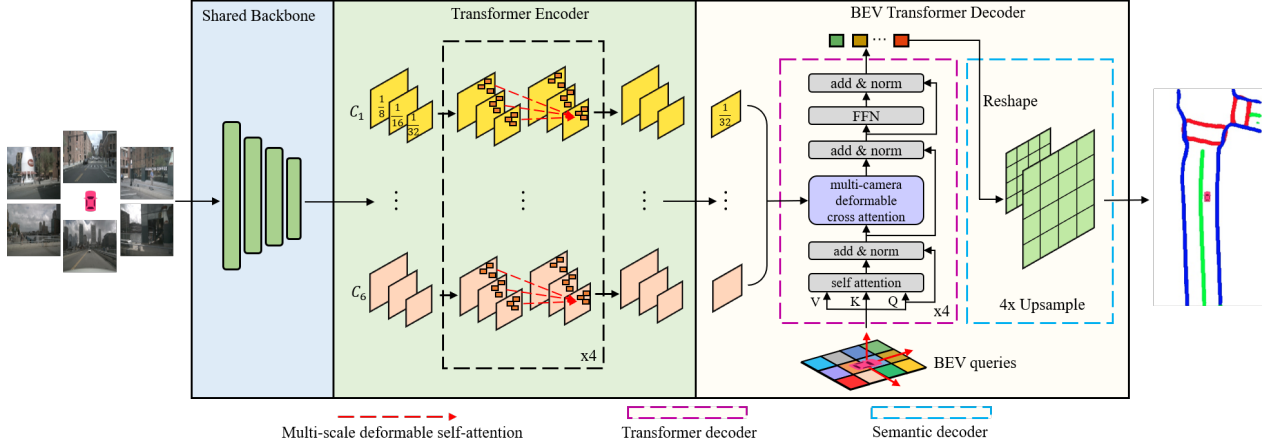


Figure 3. Overview of the proposed network. Our method consists of a Shared Backbone, a Transformer Encoder and a BEV Transformer Decoder module.

Transformer-based Semantic Segmentation. Transformer, first applied in the field of natural language processing, has been widely used in many other computer vision tasks. For example, ViT [7] introduces transformer encoder for image classification task. DETR [3] and its variant [36] are proposed for object detection by a transformer encoder-decoder architecture. SETR [35] is the first to expand transformer encoder for segmentation task. SegFormer [31] proposes an efficient semantic segmentation framework, which combines a hierarchical transformer encoder with a lightweight MLP decoder. FPN [30] introduces a full transformer encoder and decoder network for image segmentation. Adapted from DETR [3], MaskFormer [5] employs a transformer decoder to compute a set of pairs, each consisting of a class prediction and a mask embedding vector which is then combined with pixel-wise embedding from a FCN. It offers both semantic and instance segmentation in a unified manner.

Recently, transformer has been used to perform view transformation. For example, NEAT [6] uses transformer encoder and converts image features to BEV space with a MLP-based attention by traversing through all the grid in the BEV space. PYVA [32] generates BEV features from a CNN encoder-decoder structure by MLP. The feature is then enhanced by a transformer cross-attention module. Method in [25] introduces an encoder-decoder transformer block to column-wisely translate spatial features from the image to BEV. BEVFormer [13] transfers a set of 3D points back to image space and uses deformable attention to build a dense BEV map from image features. The projection is computed by camera extrinsic parameters. BEVerse [33] takes method in Lift-Splat-Shoot [21] which performs view transformation by depth estimation with camera geometry. PETR [15, 16] brings a new perspective of view transformation which encodes 3D coordinate information into 2D im-

age features. The transformer is then applied to connect the BEV space and image space by the 3D information boosted image features. Temporal information and multiple downstream tasks are also investigated in [13, 16, 33]. In this paper, we mainly focus on spatial segmentation task in BEV space. Our method has the potential to extend to associate temporal information and apply to other downstream perception tasks.

Different with these works, we introduce a BEV transformer decoder module to parse BEV semantic segmentation from the image features. The paradigm of view transformation between BEV space and image space does not require camera extrinsic parameters and it naturally supports arbitrary camera settings. Specifically, an efficient multi-camera deformable cross attention unit is designed to carry out the BEV-to-image view transformation. The queries are reshaped according to the layout of grids in the BEV, and upsampled to produce the semantic segmentation result in a supervised manner.

3. Method

Fig. 3 shows the overview of our BEVSegFormer method. It consists of three parts: (1) a Shared Backbone to process arbitrary cameras and output feature maps; (2) a Transformer Encoder to enhance the feature representation and (3) a BEV Transformer Decoder to process BEV queries by cross-attention mechanism and then parses the output queries to BEV semantic segmentation.

3.1. Shared Backbone

For a single input image, the backbone takes the input and outputs multi-scale feature maps. For multiple camera configurations, these multiple images share the same backbone and output corresponding feature maps. We take ResNet as the backbone in the experiments.

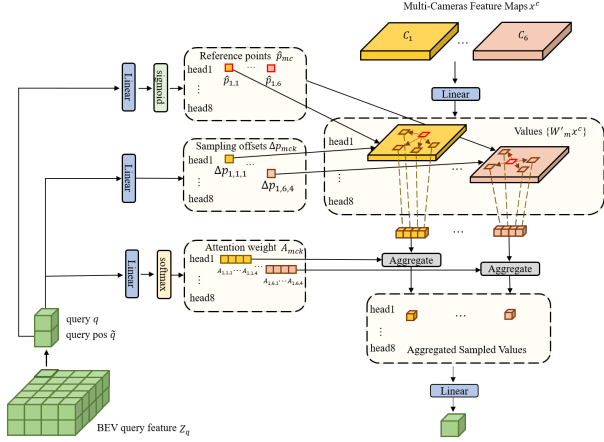


Figure 4. Illustration of the Multi-Camera Deformable Cross-Attention module in BEV Transformer Decoder.

3.2. Transformer Encoder

In the transformer encoder, we first apply 1×1 conv operators on $c3$, $c4$, $c5$ stage features from the shared backbone to get multi-scale features $\{\mathbf{x}^l\}_{l=1}^L$, where $\mathbf{x}^l \in \mathbb{R}^{C_l \times H_l \times W_l}$, C_l is channel number of feature map, H_l and W_l denotes the height and width of the feature map at l -th scale. Similar to [36], we apply multi-scale deformable self attention module separately on feature maps generated by each scale, and create additional learnable scale level position embedding for each scale feature map. Multi-scale deformable self attention module does not require to compute the dense attention map and only focuses on a set of sampling points near a reference point. The transformer encoder outputs an enhanced multi-scale features for each camera.

3.3. BEV Transformer Decoder

Our BEV transformer decoder includes a transformer decoder to compute cross-attention between BEV queries and multi-camera feature maps, and a semantic decoder to parse queries to BEV segmentation result.

In transformer decoder, we construct $H_q \times W_q$ queries on 2D BEV space and these BEV queries are then treated as $N_q = H_q \times W_q$ regular queries in the cross-attention module. The dense BEV query embeddings are denoted as $\mathbf{z}_q \in \mathbb{R}^{C \times N_q}$. We only use the smallest resolution (1/32 of original input resolution) of the multi-scale feature maps as input of the transformer decoder.

We adapt deformable attention module in deformable DETR [36] to a multi-camera deformable cross-attention module which is able to transform feature maps from multiple cameras to BEV queries, without the need of camera intrinsic and extrinsic parameters. Mathematically, let q be a query element in \mathbf{z}_q and its reference point \hat{p}_q , the multi-camera deformable attention module is written as:

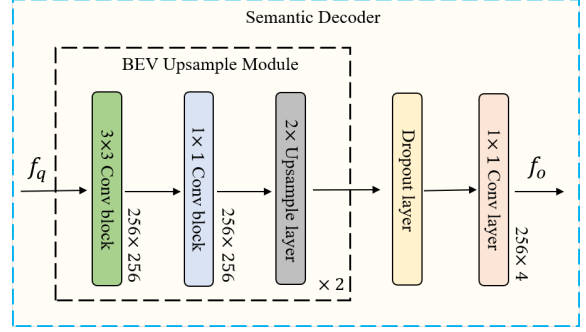


Figure 5. Illustration of Semantic Decoder.

$$\text{MultiCameraDeformAttn}(\mathbf{z}_q, \hat{p}_q, \{\mathbf{x}^c\}_{c=1}^{N_c}) = \sum_{m=1}^M \mathbf{w}_m \left[\sum_{c=1}^{N_c} \sum_{k=1}^K A_{mcqk} \cdot \mathbf{w}'_m \mathbf{x}^c(\phi_c(\hat{p}_q) + \Delta \mathbf{P}_{mcqk}) \right], \quad (1)$$

where (m, c, k) indicates the index of the attention head, camera, and sampling point, respectively. $\mathbf{w}_m \in \mathbb{R}^{C_v \times C}$ and $\mathbf{w}'_m \in \mathbb{R}^{C \times C_v}$ are the learnable parameter matrices of the linear projection layer, $C = C_v \times M$ by default. $\Delta \mathbf{P}_{mcqk}$ and A_{mcqk} are the sampling offset and attention weight of the k -th sampling point in the c -th camera and m -th attention head. The scalar attention weight A_{mcqk} is normalized to sum as 1. $\phi_c(\hat{p}_q)$ re-scales the normalized coordinates \hat{p}_q to the input feature map.

Fig. 4 shows the overall structure of the multi-camera deformable attention. For each BEV query q , we apply a learnable linear projection layer on its position embedding to obtain the 2D coordinate of reference points $\hat{p}_q \in \mathbb{R}^{M \times N_c \times 2}$, and then use the sigmoid function to normalize these coordinates. Two learnable linear projection layers are used to predict the offset of the sampling points relative to the reference point, and the attention weight of these sampling points. Finally, the camera features at the sampled positions are aggregated by the attention weights to generate a new query. Different from shared reference points in multi-scale features in Deformable DETR [36], we learn independent reference points on multi-camera feature maps, so that the network can automatically select different locations of reference points on multi-camera features.

3.4. BEV Semantic Decoder

Fig. 5 illustrates the BEV semantic decoder architecture. In semantic decoder, we reshape the BEV query features $\mathbf{z}_q \in \mathbb{R}^{C \times N_q}$ from the transformer decoder into a 2D spatial feature $\mathbf{f}_q \in \mathbb{R}^{C \times H_q \times W_q}$. The 2D spatial feature \mathbf{f}_q is then fed to a two-stage BEV Upsample Module, a dropout layer and a 1×1 convolution layer to compute the final output $\mathbf{f}_o \in \mathbb{R}^{C_{seg} \times H_q \times W_q}$. Each stage of the BEV Upsample Module consists of a 3×3 convolution block, a 1×1 convolution block and a $2 \times$ bilinear interpolation operation. A convolu-

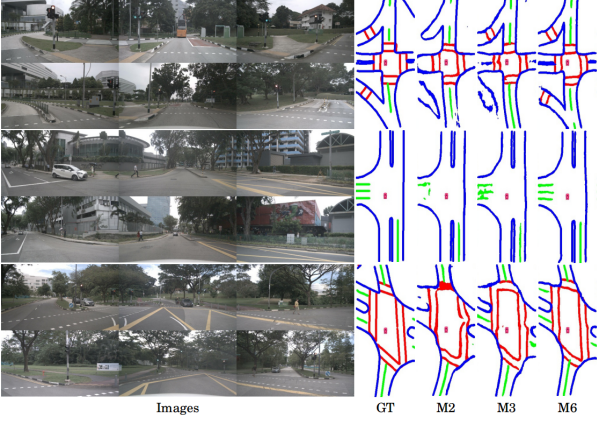


Figure 6. Examples of BEV segmentation results of surrounding cameras on nuScenes val set.

tional block includes a convolution layer, a BN layer and a ReLU activation function. The dropout layer with a dropout rate of 0.1 is applied to the feature map after restoring the resolution.

4. Experiment

4.1. Dataset

nuScenes Dataset. nuScenes [1] dataset is a large-scale autonomous driving dataset which consists of 1,000 sequences from 6 surrounding cameras (front left, front, front right, back left, back, back right). In total, it has a training set with 28,130 images and a validation set of 6,019 images. Three classes, lane dividers, lane boundaries and pedestrian crossings are available to evaluate the BEV segmentation. We use all the surrounding cameras and the front camera in the experiments.

Nullmax Front Camera Dataset. We collect a dataset from Shanghai highway with a front camera equipped. The dataset includes various scenarios such as crowd traffic, on and off ramp, shadow, lane changing and cut-in. The dataset is divided into 3,905 images for training and 976 images for validation. Traffic lanes are annotated for the evaluation.

4.2. Experimental Settings

We conduct experiments on the nuScenes dataset with the same setting of HDMapNet [12]. Ego vehicle localization on HD-map is used to define the region of BEV. By using surrounding cameras, the BEV is set to $[-30\text{m}, 30\text{m}] \times [-15\text{m}, 15\text{m}]$ around the ego vehicle. Only with the front camera enabled, the BEV area is set to $[0\text{m}, 60\text{m}] \times [-15\text{m}, 15\text{m}]$. Road structure is represented as line segments with 5-pixel width. The groundtruth mask is set to 400×200 . Following STSU [2], we use images of 448×800 as the input of the network. Similarly, the BEV area in Nullmax front camera dataset is set to $[0\text{m}, 80\text{m}] \times [-10\text{m}, 10\text{m}]$. The groundtruth mask is of size 512×128 . Traffic lane is

3-pixel width in the groundtruth mask. The input image size is 384×640 on the Nullmax dataset.

We follow deformable DETR [36] for the network design. A weighted cross entropy loss of [1, 15, 15, 15] is used in the experiments. $M = 8$ and $K = 16$ are set for multi-camera deformable attention of BEV transformer decoder. The embedding dimension in all transformer modules is set to 256, and the feature dimension of FFN module is set to 512. We apply data augmentation including random horizontal flips, random brightness, random contrast, random hue and random swap channels. The network is optimized by AdamW [17] optimizer with a weight decay of 10^{-4} . The initial learning rate of the backbone and the transformer is set to 10^{-5} , 10^{-4} and decreased to 10^{-6} and 10^{-5} at 100-th epoch. We train models on 4 RTX 3090 GPUs and the per-GPU batch size is 1. All the models are trained from scratch with 120 epochs.

4.3. Evaluation Metrics and Results

For quantitative evaluation, we measure the intersection-over-union (IoU) between the segmentation results and the groundtruth. IoU and mIoU are calculated by

$$\text{IoU}(S_p, S_g) = \frac{|(S_p \cap S_g)|}{|(S_p \cup S_g)|}, \quad (2)$$

$$\text{mIoU}(S_p, S_g) = \frac{1}{N} \sum_{n=1}^N \text{IoU}_n(S_p, S_g), \quad (3)$$

where $S_p \in \mathbb{R}^{H_g \times W_g \times N}$ and $S_g \in \mathbb{R}^{H_g \times W_g \times N}$ are the predicted results and the groundtruth. H_g and W_g denotes the height and width of the groundtruth mask and N is the number of classes in the dataset.

We compare our method with previous state-of-the-art methods on the nuScenes dataset. All these methods do not utilize temporal information for BEV semantic segmentation. The results are summarized in Table. 1. Experimental results show that our method outperforms the previous state-of-the-art approaches on the nuScenes val set. It surpasses the HDMapNet (surr) [12] by a large margin of +10.48 Divider IoU, +13.89 Ped Crossing IoU, +10.47

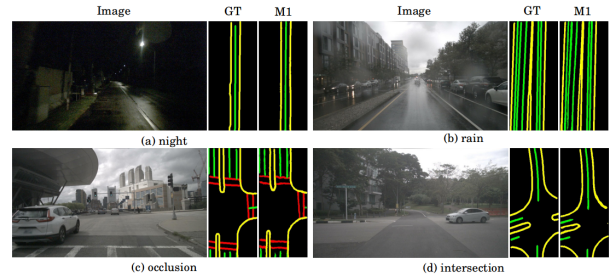


Figure 7. Examples of BEV segmentation results of front camera on nuScenes val set.

Table 1. Comparison to state of the art methods on nuScenes dataset without using temporal information. ‘†’ are the results reported in HDMaNet [12].

Method	IoU			
	Divider	Ped Crossing	Boundary	All
IPM (B)†	25.5	12.1	27.1	21.6
IPM (BC)†	38.6	19.3	39.3	32.4
Lift-Splat-Shoot† [21]	38.3	14.9	39.3	30.8
VPN† [20]	36.5	15.8	35.6	29.3
HDMaNet (Surr)† [12]	40.6	18.7	39.5	32.9
BEVSegFormer	51.08	32.59	49.97	44.55

Table 2. Result of BEV segmentation of front camera on nuScenes dataset.

#Queries	IoU			
	Divider	Ped Crossing	Boundary	All
1250	34.56	15.41	30.07	26.68
5000	37.54	17.82	34.25	29.87

Table 3. Result of BEV segmentation on Nullmax Front Camera dataset.

Method	IoU		
	Background	Lane	All
M1	92.88	67.04	79.96

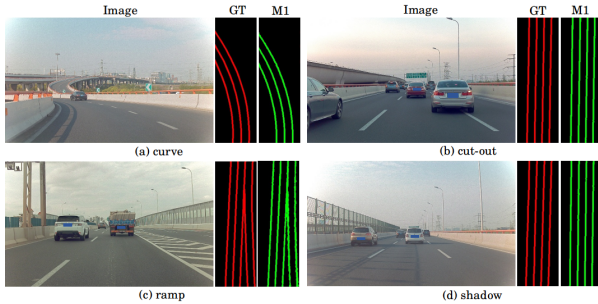


Figure 8. Examples of BEV segmentation results on Nullmax Front Camera val set.

Boundary IoU and +11.65 All Classes IoU. Segmentation examples are shown in Fig. 6.

To investigate that our method has the ability to process arbitrary cameras. We report results only using the front camera on nuScenes dataset in Table. 2. Some segmentation results are shown in Fig. 7. Only using a front camera, our method also obtains competitive results.

Similar observation of BEV segmentation on Nullmax front camera dataset is listed in Table. 3. Fig. 8 provides examples of challenging driving scenarios (curve, cut-out, ramp and shadow).

We also explore the impact of the number of queries in the BEV space. The model with 5,000 BEV queries (1/4 of

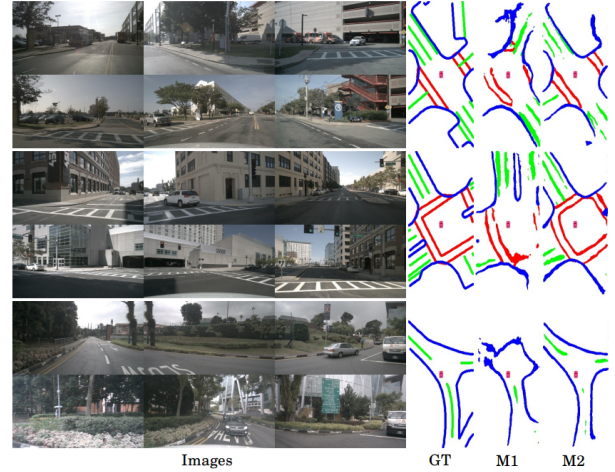


Figure 9. Results of the ablation study analyzing the effects of camera position embedding on nuScenes val set.

groundtruth mask) achieves better results compared to the model with 1,250 queries. Throughout the experiments, the query number is set to 5,000.

4.4. Ablations

In this section, we analyze the effects of the proposed components via an ablation study. We study the standard multi-head self attention module, multi-scale deformable self attention module, standard multi-head cross attention module, multi-camera deformable cross attention module and number of layers in the encoder and decoder of our method. Results of various backbones are also investigated. Additionally, the effect of camera position embedding in multi-camera attention is provided. The experimental settings and results are given in Table. 4. The methods with different protocols are denoted as M1 to M7.

Learnable Camera Position Embedding.

We find that a learnable camera position embedding is important for the standard cross attention module in multi-camera decoder, while it does not provide a positive effect of multi-camera deformable cross attention module. In Table. 4, having camera position embedding with standard cross attention module, method M2 improves the segmen-

Table 4. Results of ablation study for the components of BEVSegFormer on the nuScenes val set. ‘Enc’ denotes encoder of model, ‘Dec’ denotes decoder of model, ‘S’ denotes standard multi-head self attention module and standard multi-head cross attention module, ‘D’ denotes multi-scale deformable self attention module and multi-camera deformable cross attention module. ‘#Enc’ and ‘#Dec’ denotes number of encoder block and decoder block. ‘CE’ denotes camera position embedding.

Method	Backbone	Enc	Dec	#Enc	#Dec	CE	IoU			
							Divider	Ped Crossing	Boundary	All
M1	ResNet-34	S	S	2	2		44.23 (+0.00)	24.32 (+0.00)	42.25 (+0.00)	36.93 (+0.00)
M2	ResNet-34	S	S	2	2	✓	45.43 (+1.20)	26.57 (+2.25)	45.30 (+3.05)	39.10 (+2.17)
M3	ResNet-34	S	D	2	2		47.84 (+3.61)	28.88 (+4.56)	46.74 (+4.49)	41.15 (+4.22)
M4	ResNet-34	D	D	2	2		49.03 (+4.80)	30.55 (+6.23)	48.05 (+5.80)	42.54 (+5.61)
M5	ResNet-34	D	D	4	4		50.79 (+6.56)	32.39 (+8.07)	49.84 (+7.59)	44.34 (+7.41)
M6	ResNet-101	D	D	4	4		51.08 (+6.85)	32.59 (+8.27)	49.97 (+7.72)	44.55 (+7.62)
M7	ResNet-101	D	D	4	4	✓	50.29 (+6.06)	31.82 (+7.50)	50.18 (+7.93)	44.10 (+7.17)

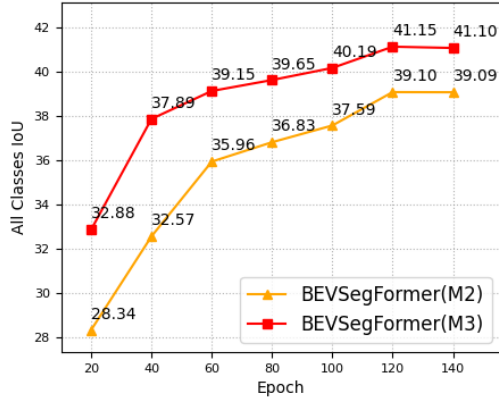


Figure 10. Convergence curves of BEVSegFormer (M2) and BEVSegFormer (M3) on nuScenes val set. The same training schedules is used in both models.

tation results of M1. Results in Fig. 9 demonstrate that method incorporating camera position embedding obtains superior results at scenarios that the ego vehicle makes a turn at the intersections. However, by using multi-camera deformable cross attention module, M7 with a learnable camera position embedding does not boost semantic segmentation results of M6.

We argue that in standard multi-camera cross attention module, the position embedding of image features is relative to each camera itself and does not have ability to discriminate different cameras. Therefore, the spatial information of the camera is critical for semantic segmentation tasks using standard multi-camera cross attention module. In our proposed multi-camera deformable cross attention module, a BEV query learns a separate reference point for each camera. It automatically encode camera position information into the cross attention step. It turns out that an additional and redundant camera position embedding provides marginal effect for the final semantic segmentation accuracy.

Multi-camera Transformer Decoder. Experiments in Table. 4 show that compared to methods M1 and M2 using standard multi-head cross attention in the decoder, method M3 with multi-camera deformable cross attention in the decoder obtains significant performance gains. Multi-camera deformable decoder does not require to compute the dense attention over feature maps from all the cameras. In this way, the learned decoder attention has strong correlation with the BEV segmentation results, which also accelerates the training convergence. Fig. 10 shows the convergence curves of the BEVSegFormer (M3) and BEVSegFormer (M2) models on the nuScenes val set. The former only needs half the training epochs ($epoch = 60$) to obtain a higher IoU of all classes than the latter ($epoch = 120$).

In order to investigate the multi-camera deformable cross-attention, we analyze the attention map in Fig. 11. Specifically, given a query of road structure on the BEV space, we visualize the corresponding sampling points and attention weights in the last layer of the multi-camera deformable cross-attention. It shows that the sampling points of large weights are clustered near the feature maps of the most relative camera. For example, in the top left of Fig. 11, the query (black point) on BEV space is observed in the front camera, its sampling points with large cross-attention weights appear in the feature map of the front camera. Even though in some cases that the sampling points are not precisely linked with the desired query, the image features from the relative camera have enough receptive fields to cover the input query on BEV space.

Multi-scale Transformer Encoder. In Table. 4, We also compare the segmentation results using standard transformer encoder (M3) and deformable encoder (M4) which utilizes multi-scale images. Similar to deformable DETR [36], the deformable encoder module not only accelerates the model convergence, but also provides multi-level image features which in turn to boost the BEV segmentation results.

Number of Encoder Block and Decoder Block. Ta-



Figure 11. Visualization of Multi-Camera Deformable Cross-Attention weight map in BEV Transformer Decoder module. We visualized the attention weight map of the query corresponding to lane dividers (in green), pedestrian crossings (in red), lane boundaries (in blue) and camera overlap regions on the BEV. Redder pixels indicate larger values of attention weights.

ble. 4 shows the segmentation results using different numbers of encoder layers and decoder layers. Compared to method M4, the segmentation results of method M5, which has more encoder and decoder layers, are further boosted.

Shared Backbone. Lastly, we study the effect of the backbones. Results of method using ResNet-101 slightly improves the results of method with ResNet-34 as backbone.

5. CONCLUSIONS

In this paper, we propose a deformable transformers based method for BEV semantic segmentation from arbitrary camera rigs. Our method adapts the recent detection method, deformable DETR, which processes a single image to BEV semantic segmentation of single or multiple images. We show that our transformer based method outperforms previously reported state-of-the-art approaches. By making use of the deformable BEV decoder, a novel component to parse BEV semantic results from the encoded image features, our method achieves additional boost of the performance. In future work, we would like to explore memory-efficient and more interpretable transformer-based methods for BEV semantic segmentation in autonomous driving.

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [2] Yigit Baran Can, Alexander Liniger, Danda Pani Paudel, and Luc Van Gool. Structured bird’s-eye-view traffic scene understanding from onboard images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15661–15670, 2021.
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020.
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [5] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [6] Kashyap Chitta, Aditya Prakash, and Andreas Geiger. Neat: Neural attention fields for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15793–15803, 2021.
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [8] Divam Gupta, Wei Pu, Trenton Tabor, and Jeff Schneider. Sbevnet: End-to-end deep stereo layout estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 523–532, 2022.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Nouredin Hendy, Cooper Sloan, Feng Tian, Pengfei Duan, Nick Charchut, Yuesong Xie, Chuang Wang, and James Philbin. Fishing net: Future inference of semantic heatmaps in grids. *arXiv preprint arXiv:2006.09917*, 2020.

- [11] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15273–15282, 2021.
- [12] Qi Li, Yue Wang, Yilun Wang, and Hang Zhao. Hdmapnet: A local semantic map learning and evaluation framework. *arXiv preprint arXiv:2107.06307*, 2021.
- [13] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022.
- [14] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017.
- [15] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. *arXiv preprint arXiv:2203.05625*, 2022.
- [16] Yingfei Liu, Junjie Yan, Fan Jia, Shuailin Li, Qi Gao, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petrv2: A unified framework for 3d perception from multi-camera images. *arXiv preprint arXiv:2206.01256*, 2022.
- [17] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [18] Kaustubh Mani, Swapnil Daga, Shubhika Garg, Sai Shankar Narasimhan, Madhava Krishna, and Krishna Murthy Jatavallabhula. Monolayout: Amodal scene layout from a single image. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1689–1697, 2020.
- [19] Mong H Ng, Kaahan Radia, Jianfei Chen, Dequan Wang, Ionel Gog, and Joseph E Gonzalez. Bev-seg: Bird’s eye view semantic segmentation using geometry and semantic point cloud. *arXiv preprint arXiv:2006.11436*, 2020.
- [20] Bowen Pan, Jiansun Sun, Ho Yin Tiga Leung, Alex Andonian, and Bolei Zhou. Cross-view semantic segmentation for sensing surroundings. *IEEE Robotics and Automation Letters*, 5(3):4867–4873, 2020.
- [21] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision*, pages 194–210. Springer, 2020.
- [22] Lennart Reiher, Bastian Lampe, and Lutz Eckstein. A sim2real deep learning approach for the transformation of images from multiple vehicle-mounted cameras to a semantically segmented image in bird’s eye view. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7. IEEE, 2020.
- [23] Thomas Roddick and Roberto Cipolla. Predicting semantic map representations from images using pyramid occupancy networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11138–11147, 2020.
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [25] Avishkar Saha, Oscar Mendez Maldonado, Chris Russell, and Richard Bowden. Translating images into maps. *arXiv preprint arXiv:2110.00966*, 2021.
- [26] Samuel Schulter, Menghua Zhai, Nathan Jacobs, and Manmohan Chandraker. Learning to look around objects for top-view representations of outdoor scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 787–802, 2018.
- [27] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):640–651, 2016.
- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [29] Yue Wang, Vitor Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. *arXiv preprint arXiv:2110.06922*, 2021.
- [30] Sitong Wu, Tianyi Wu, Fangjian Lin, Shengwei Tian, and Guodong Guo. Fully transformer networks for semantic image segmentation. *arXiv preprint arXiv:2106.04108*, 2021.
- [31] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34, 2021.
- [32] Weixiang Yang, Qi Li, Wenxi Liu, Yuanlong Yu, Yuexin Ma, Shengfeng He, and Jia Pan. Projecting your view attentively: Monocular road scene layout estimation via cross-view transformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15536–15545, 2021.
- [33] Yunpeng Zhang, Zheng Zhu, Wenzhao Zheng, Junjie Huang, Guan Huang, Jie Zhou, and Jiwen Lu. Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving. *arXiv preprint arXiv:2205.09743*, 2022.
- [34] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [35] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021.
- [36] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations*, 2020.