# FSD-BEV: Foreground Self-Distillation for Multi-view 3D Object Detection

Zheng Jiang<sup>1,3,4</sup>\*<sup>©</sup>, Jinqing Zhang<sup>2</sup>\*<sup>©</sup>, Yanan Zhang<sup>2</sup><sup>©</sup>, Qingjie Liu<sup>1,2,5</sup><sup>†</sup><sup>©</sup>, Zhenghui Hu<sup>1,2</sup><sup>†</sup><sup>©</sup>, Baohui Wang<sup>3</sup>, and Yunhong Wang<sup>1,2</sup><sup>©</sup>

<sup>1</sup> Hangzhou Innovation Institute, Beihang University, Hangzhou, China

<sup>2</sup> State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China

<sup>3</sup> School of Software, Beihang University, Beijing, China

<sup>4</sup> Shanghai ZEEKR Blue New Energy Technology Co., Ltd. <sup>5</sup> Zhongguancun Laboratory, Beijing, China

Abstract. Although multi-view 3D object detection based on the Bird's-Eye-View (BEV) paradigm has garnered widespread attention as an economical and deployment-friendly perception solution for autonomous driving, there is still a performance gap compared to LiDAR-based methods. In recent years, several cross-modal distillation methods have been proposed to transfer beneficial information from teacher models to student models, with the aim of enhancing performance. However, these methods face challenges due to discrepancies in feature distribution originating from different data modalities and network structures, making knowledge transfer exceptionally challenging. In this paper, we propose a Foreground Self-Distillation (FSD) scheme that effectively avoids the issue of distribution discrepancies, maintaining remarkable distillation effects without the need for pre-trained teacher models or cumbersome distillation strategies. Additionally, we design two Point Cloud Intensification (PCI) strategies to compensate for the sparsity of point clouds by frame combination and pseudo point assignment. Finally, we develop a Multi-Scale Foreground Enhancement (MSFE) module to extract and fuse multi-scale foreground features by predicted elliptical Gaussian heatmap, further improving the model's performance. We integrate all the above innovations into a unified framework named FSD-BEV. Extensive experiments on the nuScenes dataset exhibit that FSD-BEV achieves state-of-the-art performance, highlighting its effectiveness. The code and models are available at: https://github.com/CocoBoom/fsd-bev.

Keywords: BEV · Self-Distillation · Multi-view 3D Object Detection

# 1 Introduction

3D object detection is a fundamental task in autonomous driving. Although LiDAR-based methods [31, 34, 35, 38] have demonstrated excellent performance,

 $<sup>\</sup>star$  Equal contribution.  $\dagger$  Corresponding author.



**Fig. 1:** Comparison of cross-modal distillation framework and our self-distillation framework. The hard labels denote the depth maps and foreground segmentation generated by LiDAR point clouds, the soft labels are those predicted by the student. The teacher branch and student branch share the same image features and thus mitigate the distribution discrepancies between the distillation targets.

their limitations, such as the absence of color and texture information and the high hardware costs involved, have sparked a growing interest in approaches utilizing multi-view cameras. Early multi-view camera-based methods [20, 24, 25] initially performed object detection tasks independently on individual images, followed by cross-view integration. However, this approach often resulted in suboptimal performance and hindered the coherence of subsequent tasks. Subsequent research has shifted focus towards the Bird's-Eye-View (BEV) paradigm, offering a stronger universal intermediate representation for various views and modalities, and facilitating the integration of temporal information.

Although existing BEV-based multi-view 3D object detection methods have achieved competitive precision, there are still certain performance gaps compared to methods based on LiDAR. Distillation methods are employed to mitigate the performance gap between different modalities. Most of the BEV-based distillation methods [2,9,28,36] fall into the framework illustrated in Fig. 1a. A pre-trained teacher model transforms the LiDAR point clouds or multi-modal inputs into the frozen teacher BEV features, which serve as the inflexible guidance for the BEV features generated by the student. However, the distribution discrepancies between the distillation targets are the shared challenge encountered by these distillation methods, which are caused by differences in data modality, network structures and other factors. It makes it necessary to adopt cumbersome distillation strategies [9, 28, 36] to ensure the effectiveness.

In this paper, we propose the foreground self-distillation scheme for the BEVbased 3D objects, aiming to enhance model performance through the collaborative growth of teacher and student branches. Unlike the previous cross-modal distillation methods, a separate teacher model is not required in our framework; instead, a teacher branch is attached to the student model. As illustrated in Fig. 1b, the teacher branch utilizes the hard labels generated by LiDAR point clouds to obtain high-quality teacher BEV features and offer guidance to the student branch, while conversely, the student compensates the teacher with the predicted soft labels to fill the vacancy of the hard labels. It is worth noting that the teacher and student BEV features share the same source of image features in our self-distillation framework. It avoids severe distribution discrepancies and allows joint progress of the student and the teacher, which raises the upper bound on the distillation effect. Another challenge when distilling BEV features is that a large ratio of background provides trivial benefits and even disturbs the alignment in the foreground. Other methods commonly design complicated strategies to pay more attention to the foreground during the distillation process. In our scheme, the same purpose can be easily achieved by filtering out the background in both teacher and student BEV features following the way of SA-BEV [32]. We name the whole framework described above Foreground Self-Distillation (FSD).

The performance of the teacher branch in FSD heavily depends on the quality of the hard labels generated by point clouds. As a result, we design two Point Cloud Intensification (PCI) strategies to address the sparsity of point cloud data, namely, merging frame information and assigning pseudo points to objects that have no associated points. The sparsity of the hard labels is well mitigated, providing better guidance for feature synthesis. Furthermore, we develop a Multi-Scale Foreground Enhancement (MSFE) module to extract and fuse multi-scale foreground features by predicted elliptical Gaussian heatmap. This module effectively helps the student branch to predict more precise soft labels and thus improve the performance of the whole framework.

The strengths of our foreground self-distillation approach are: (1) No need for additional auxiliary models and pre-training processes; (2) Without cumbersome distillation strategies; (3) Avoidance of severe distribution discrepancies between the distillation targets; (4) Joint progress of the student and teacher model. We integrate the above designs into one framework named FSD-BEV. Extensive experiments validate the effectiveness of our method, demonstrating outstanding performance. In summary, the major contributions of this paper are:

- We propose the first Foreground Self-Distillation framework for BEV-based multi-view 3D object detection, effectively mitigating the distribution discrepancies of the distilled features and background noise interference inherent in cross-modal distillation.
- We design two Point Cloud Intensification strategies to compensate for the sparsity of point clouds by frame combination and pseudo point assignment, further intensifying the performance ceiling of the teacher model.
- We develop a Multi-Scale Foreground Enhancement module to extract and fuse multi-scale foreground features by predicted elliptical Gaussian heatmap, further improving the model's performance.

# 2 Related Work

# 2.1 Multi-view 3D Object Detection

Multi-view 3D object detection is a critical task for autonomous driving systems, aiming to harness information from various perspectives. The Bird's-Eye View (BEV) representation, offering a unified representation of multiple views, divides research in this field into two paradigms: attention-based and depth-based. DETR3D [26] is a representative of the attention-based paradigm, utilizing sparse 3D queries to index 2D features extracted from images. BEV-Former [13] introduced temporal sequences and incorporated spatiotemporal attention mechanisms. BEVFormerV2 [30] extended its focus to the gains from 2D tasks. The PETR series [17, 18, 23] introduced 3D positional information to queries and underwent various improvements.

In another branch, BEVDet [8] is the pioneer that combines LSS [22] with the detection head, treating depth information as an explicit prediction for detection tasks. BEVDepth [12] uses LiDAR ground truth as depth supervision, significantly enhancing detection performance. BEVStereo [11] further refines depth predictions by utilizing stereo vision matching. STS [27] enhances depth quality by leveraging geometric correspondences between adjacent frames. However, regardless of the type of method, existing multi-view 3D object detection still lags behind LiDAR-based methods due to insufficiently acquired geometric information. In this paper, we focus on designing more suitable distillation techniques to narrow the performance gap between multi-view 3D object detection and LiDAR-based methods.

#### 2.2 Knowledge Distillation

Knowledge distillation is employed for capacity learning between two models. It enhances the performance of the model without altering the student model structure and has been widely applied [3, 4, 6]. Self-distillation [29, 33] is another approach that implements distillation between features at different levels or transferring knowledge across various periods. In the field of 3D object detection, LIGA-stereo [5] utilizes LiDAR-based detectors to guide Stereo-based detectors. UVTR [10] unifies multi-modal representations in voxel space, enabling accurate single-modal or cross-modal 3D detection. BEVDistill [2] performs cross-modal distillation in a unified BEV space. UniDistill [36] proposes three distillation approaches in three stages. DistillBEV [28] proposes a balanced design for the fusion of multi-scale and temporal information. However, existing methods generally require additional training of teacher models and suffer from the divergence in modality and model structure between teacher and student, resulting in limited performance gains from knowledge distillation. In this paper, we propose a foreground self-distillation scheme that effectively alleviates distribution discrepancies, achieving excellent distillation effects without additional pre-trained models.

# 3 Method

In this section, we will present the details of FSD-BEV, the first self-distillation approach for the BEV-based 3D detectors. The overall architecture of FSD-BEV is illustrated in Fig. 2. Unlike the former BEV-based cross-modal distillation methods that employ extra pre-trained teacher models and use the cumbersome



Fig. 2: The overall architecture of the proposed FSD-BEV. The features enhanced by the foreground heatmap are fed into the View Transformation Module to generate the student BEV. The teacher branch generates the teacher BEV by combining hard labels with soft labels from the student branch. Subsequently, they are concatenated along the batch dimension for subsequent joint training and undergo distillation operations before entering the detection head.

distillation strategy, Foreground Self-Distillation achieves significant accuracy improvement by simply doing feature alignment within a single model. Subsequently, we conduct several Point Cloud Intensification procedures to build more truthful teacher BEV features and further improve the effect of distillation. Finally, we effectively aggregate foreground information from different scale feature maps using the Multi-Scale Foreground Enhancement strategy, extracting more useful 2D image features.

# 3.1 Foreground Self-Distillation

The former BEV-based cross-modal distillation methods [2, 9, 28, 36] generally require a heavier LiDAR-based or multi-modal teacher model to serve as the guidance. And BEV features are commonly employed by these methods as the distillation medium. However, the distribution gap between the BEV features of teacher and student features poses challenges in cross-modal distillation. Such a distribution gap may arise from differences in data modality, network structures, and other factors. The focus of distillation methods lies in transferring features effectively under distribution disparities, necessitating the design of clever strategies to maximize distillation efficiency. Since the imitation in the background region offers trivial benefits to the detection accuracy, it is a natural idea to



Fig. 3: Details of BEV feature generation. Due to the sparse nature of the depth label, the teacher branch cannot acquire dense foreground information. To address this limitation, we use soft labels generated by the student branch to assist the teacher branch in generating denser depth maps.

select foreground regions by projecting ground truths boxes onto the BEV and perform delicate distillation operations there.

In the proposed Foreground Self-Distillation, significant performance improvement is achieved without the need for additional auxiliary models or elaborate distillation strategies. Unlike employing pre-trained teacher models in conventional distillation approaches, the teacher and student engage in a joint learning process in our method. While the teacher and student BEV features share the same context feature resources, more accurate depth and semantic information help construct the high-performing teacher BEV features, providing ongoing guidance for the student BEV features. The foreground segmentation is also introduced to generate BEV features that only contain foreground information, which abandons the useless imitation of the background region and avoids noise interference. The foreground segmentation also dramatically increases the accuracy of the teacher branch, consequently strengthening the distillation efficiency.

**Student BEV** As illustrated in Fig. 3, we generate the student BEV features in the View Transformation Module. Following conventional procedures [12], we predict context features  $\mathbf{C}$  and predicted depth maps  $\mathbf{D}$  from the input image feature map. To filter out the background in the BEV features, the predicted foreground segmentation  $\mathbf{S}$  is required. It can be obtained by reusing the Depth-Net designed for predicting  $\mathbf{D}$  or by the specialized module described in Sec. 3.3. We leverage SA-BEVPool [32] to generate BEV features that contain only foreground information. The student BEV features can be represented by:

$$\mathbf{B}_s = \text{SA-BEVPool}(\mathbf{C}, \mathbf{D}, \mathbf{S}), \tag{1}$$

which means **C**, **D** and **S** are taken as the input of SA-BEVPool and the student BEV features carrying only foreground information are created.

**Teacher BEV** In another branch of Fig. 3, we generate the teacher BEV features by injecting accurate depth and semantic information. Specifically, we utilize the ground truth depth maps  $\hat{\mathbf{D}}$  and foreground segmentation  $\hat{\mathbf{S}}$ , obtained from LiDAR point clouds as stated in SA-BEV [32], to replace the relatively unfaithful  $\mathbf{D}$  and  $\mathbf{S}$ . While naming the ground truth as the hard labels, the predicted  $\mathbf{D}$  and  $\mathbf{S}$  are known as the soft labels. Although hard labels reveal accurate scene information, the sparsity of point clouds can lead to the label missing. To address this issue, soft labels are utilized to fill the missing part of the hard labels and the combined labels can be represented by:

$$\bar{\mathbf{D}} = \mathbf{M}\hat{\mathbf{D}} + (1 - \mathbf{M})\mathbf{D}$$
(2)

$$\tilde{\mathbf{S}} = \mathbf{M}\tilde{\mathbf{S}} + (1 - \mathbf{M})\mathbf{S},\tag{3}$$

where  $\mathbf{M}$  denotes the valid mask of the hard labels. The value of  $\mathbf{M}$  is 1 where the hard labels are available and 0 otherwise. The teacher BEV features are then generated by:

$$\mathbf{B}_t = \text{SA-BEVPool}(\mathbf{C}, \mathbf{D}, \mathbf{S}). \tag{4}$$

Since both student and teacher BEV features inherit the information of  $\mathbf{C}$  in our self-distillation framework, the gap between them is much smaller than in the former cross-modal distillation framework. Meanwhile, part of the soft labels predicted by the student branch are also involved in constructing teacher BEV features, making it easier for the student to imitate the teacher.

**Collaborative Training** It is not easy to directly align  $\mathbf{B}_s$  and  $\mathbf{B}_t$ . The former distillation methods attach additional adaptation modules, which map the student BEV features to the teacher BEV features. Instead of following the conventional approach, we concatenate the  $\mathbf{B}_s$  and  $\mathbf{B}_t$  together along the batch dimension. Subsequently, they jointly undergo processing through the BEV encoder, resulting in high-level BEV features, which can be represented by:

$$\dot{\mathbf{B}}_s = \text{BEVEncoder}(\mathbf{B}_s) \tag{5}$$

$$\mathbf{B}_t = \text{BEVEncoder}(\mathbf{B}_t). \tag{6}$$

The alignment is then applied between  $\hat{\mathbf{B}}_s$  and  $\hat{\mathbf{B}}_t$ . It can be found that the BEV encoder plays the same role as the adaptation modules without increasing the parameters. The BEV encoder can also be regarded as a feature filter, which filters the incorrect information in  $\mathbf{B}_s$  and makes the filtered  $\hat{\mathbf{B}}_s$  similar to  $\hat{\mathbf{B}}_t$ .

It is noteworthy that  $\hat{\mathbf{B}}_s$  and  $\hat{\mathbf{B}}_t$  should also be concatenated and jointly fed into the detection head. If  $\hat{\mathbf{B}}_t$  are not supervised by the detection losses, it will move to  $\hat{\mathbf{B}}_s$  after the alignment, which is opposite to our purpose. When  $\hat{\mathbf{B}}_s$ and  $\hat{\mathbf{B}}_t$  are simultaneously supervised, the relative unfaithful  $\hat{\mathbf{B}}_s$  tends to incur larger detection losses than the  $\hat{\mathbf{B}}_t$ , which mean the participation of  $\hat{\mathbf{B}}_t$  will not heavily influence the training of student branch.



**Fig. 4:** Overview of Point Cloud Intensification. In (a), we enhance the current frame by merging points of static objects from the adjacent frames in time. In (b), green and red points represent foreground and background points, respectively. We employ the PPA strategy to complete missing point clouds for foreground objects.

**Distillation Loss** Benefiting from the homogeneity between  $\mathbf{B}_s$  and  $\mathbf{B}_t$  and their foreground-only characteristic, we circumvent the necessity of devising intricate distillation losses for ensuring effective information transfer. Instead, the simple  $L_2$  loss can suffice. A problem is that both  $\hat{\mathbf{B}}_s$  and  $\hat{\mathbf{B}}_t$  are not frozen during collaborative training, which means the  $L_2$  distillation loss can take a shortcut by decreasing the magnitudes of  $\hat{\mathbf{B}}_s$  and  $\hat{\mathbf{B}}_t$  instead of pulling them together. As a result, we normalize the magnitudes of  $\hat{\mathbf{B}}_s$  and  $\hat{\mathbf{B}}_t$  before calculating the distillation loss. The distillation loss can be represented by:

$$\mathcal{L}_{distill} = \frac{1}{H \times W} \sum_{i}^{H} \sum_{j}^{W} \left\| \frac{\hat{\mathbf{B}}_{t}^{ij} - \hat{\mathbf{B}}_{s}^{ij}}{\|\hat{\mathbf{B}}_{t}^{ij}\|_{2}} \right\|_{2}.$$
(7)

#### 3.2 Point Cloud Intensification

The quality of hard labels  $\hat{\mathbf{D}}$  and  $\hat{\mathbf{S}}$  generated by LiDAR point clouds determines the performance of the teacher branch, thereby influencing the distillation results. However, the sparsity of point clouds often results in plenty of distant objects having few or no points, weakening the quality of hard labels. We introduce two kinds of Point Cloud Intensification (PCI) strategies here to make hard labels carry more information about the scene as shown in Fig. 4.

**Frame Combination** The first PCI strategy named Frame Combination (FC) uses the adjacent frames in time to provide supplementary point clouds. To avoid the points of the dynamic objects introducing errors, only the points belonging

to the stationary foreground objects will be combined, such as the parked cars, the cycles without riders and the traffic cones. We transform the point clouds of the adjacent frames into the coordinate system of the current frame and use the stationary object boxes to select the desired points. the objects in the current frame will have denser points after FC as illustrated in Fig. 4a.

**Pseudo Point Assignment** After combining the adjacent point clouds, there are still some objects that do not show up on the hard labels. They could be dynamic objects that are not suitable for Frame Combination, or be quite distant that even the adjacent frames can not provide valid points. In this situation, it is a reasonable choice to assign the approximate points to these objects in the space, which is another PCI strategy called Pseudo Point Assignment (PPA).

We first project the ground truth 3D boxes onto the image to get their corresponding 2D boxes. Each 2D box can be represented by  $(x_1, y_1, x_2, y_2)$ , where  $(x_1, y_1)$  denotes the top-left point,  $(x_2, y_2)$  denotes the bottom-right point. We select the boxes that should be assigned with a pseudo point following several criteria: 1) there are no real points within the box after Frame Combination; 2) the depth of the box is in the perception range; 3) the box has good visibility (set to 3 or 4 in the nuScenes [1] dataset). The coordinate of the pseudo point in the image coordinate system can be represented by:

$$p_{pseudo} = \left(\frac{x_1 + x_2}{2}, \frac{y_1 + y_2}{2}, d_{corner}\right),\tag{8}$$

where  $d_{corner}$  represents the minimum depth among the eight corner points of the 3D box, which is close to the surface depth of the object. As shown in Fig. 4b, the teacher BEV features can capture the missing objects after PPA.

#### 3.3 Multi-Scale Foreground Enhancement

Fusing high-scale features is a straightforward way to provide finer depth maps to the view transformation module. Meanwhile, we think the foreground is more beneficial than the background in high-scale features. We introduce a Multi-Scale Foreground Enhancement (MSFE) module as shown in Fig. 2 to facilitate the integration of extensive foreground details into the model's learning paradigm.

We define  $\mathbf{F}^4$ ,  $\mathbf{F}^8$ ,  $\mathbf{F}^{16}$  as the output feature maps of the FPN [15] layer, where the superscript represents the downsampling ratio of the feature map size from the input image size. Unlike the former BEV-based methods commonly utilize  $\mathbf{F}^{16}$  as the input features for the depth network, we predict the high-scale foreground segmentation  $\mathbf{S}^4$  by using  $\mathbf{F}^4$  to provide more detailed information.

On such a high scale, the labels generated by LiDAR point clouds are fairly sparse (about 80% of the labels are vacant) and cannot offer effective supervision. Instead, we adopt the heatmaps generated by the ground truth 2D boxes on the images as the labels. Unlike CenterNet [37] drawing circular Gaussian distributions around the object centers, we represent the confidence distribution of foreground objects by elliptical Gaussians that fill the 2D boxes, and longer dimension has larger covariance. Focal Loss is utilized as the training loss.

Once  $\mathbf{S}^4$  is predicted, we gather the foreground information of multi-scale features and integrate them through several dot product and downsampling operations. The  $\mathbf{S}^4$  is first filtered by a threshold  $\beta$  to get  $\mathbf{S}_f^4$ , which can be represented by:

$$\mathbf{S}_{f}^{4} = M_{f} \cdot \mathbf{S}^{4}, \ M_{f} = \begin{cases} 1, & \mathbf{S}^{4} \ge \beta \\ 0, & \mathbf{S}^{4} < \beta \end{cases}.$$

$$\tag{9}$$

Denoting  $DS^n$  as the downsampling operation by a factor of n, the enhanced  $\mathbf{F}_{MSFE}^{16}$  can be represented by:

$$\mathbf{F}_{MSFE}^{16} = \mathbf{F}^{16} + \mathrm{DS}^2(\mathbf{F}^8 \cdot \mathrm{DS}^2(\mathbf{S}_f^4)) + \mathrm{DS}^4(\mathbf{F}^4 \cdot \mathbf{S}_f^4)$$
(10)

After feature aggregation,  $\mathbf{F}_{MSFE}^{16}$  is supplemented with more detailed information, making the prediction based on it more precise. The high-scale  $\mathbf{S}^4$  can also be utilized by the student branch to offer a finer foreground segmentation.

# 4 Experiments

#### 4.1 Experimental Settings

**Dataset and Metrics** We conducted our experiments on the nuScenes [1] dataset, which comprises 1000 sequences captured by sensors such as cameras and LiDAR. Each sequence is approximately 20 seconds long, with 700 sequences used for training, 150 for validation, and 150 for testing. We employed official evaluation metrics, including mAP, ATE, ASE, AOE, AVE, and AAE, to assess the results for center distance, translation, scale, orientation, velocity, and attributes, respectively. Additionally, the official evaluation provides a comprehensive metric known as nuScenes Detection Score (NDS).

Implementation Detail Our model design follows the classical depth-based BEV detection model architecture. We conducted experiments using 8 NVIDIA GeForce RTX 3090 GPUs. AdamW [19] is employed as our optimizer with a learning rate set to 2e-4 and kept constant throughout the training process. Additionally, we adopted mixed precision training to save more resources. In most ablation experiments, we utilized ResNet50 as the backbone network with an image input size of  $256 \times 704$ . When comparing with other state-of-the-art models, we use larger image size and adopt heavier backbones. If not stated, the model is trained for 24 epochs with the CBGS [39] strategy. Our experiments are conducted using only one past frame with an interval of 0.5s.

#### 4.2 Main Results

We compare our method with state-of-the-art approaches on the val and test sets of nuScenes [1]. As shown in Tab. 1, when we use the smaller ResNet50 as the backbone network, our method outperforms all the methods using only

**Table 1:** Comparison with the state-of-the-art methods on the nuScenes *val* set. \* Benefited from the perspective-view pre-training. <sup>†</sup> Increasing the size of BEV features from  $128 \times 128$  to  $256 \times 256$  to enhance the granularity of distillation.

| Methods                     | Backbone                    | Resolution        | Frames | $\mathrm{mAP}\uparrow$ | $\mathrm{NDS}\uparrow$ | $\mathrm{mATE}{\downarrow}$ | $\mathrm{mASE}{\downarrow}$ | $\mathrm{mAOE}{\downarrow}$ | $\mathrm{mAVE}{\downarrow}$ | mAAE↓ |
|-----------------------------|-----------------------------|-------------------|--------|------------------------|------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-------|
| BEVDet [8]                  | ResNet50                    | $256{	imes}704$   | 1      | 0.298                  | 0.379                  | 0.725                       | 0.279                       | 0.589                       | 0.860                       | 0.245 |
| BEVDet4D [7]                | ResNet50                    | $256{	imes}704$   | 2      | 0.322                  | 0.457                  | 0.703                       | 0.278                       | 0.495                       | 0.354                       | 0.206 |
| PETRv2 [18]                 | ResNet50                    | $256{	imes}704$   | 2      | 0.349                  | 0.456                  | 0.700                       | 0.275                       | 0.580                       | 0.437                       | 0.187 |
| BEVDepth [12]               | ResNet50                    | $256{	imes}704$   | 2      | 0.351                  | 0.475                  | 0.639                       | 0.267                       | 0.479                       | 0.428                       | 0.198 |
| BEVStereo [11]              | ResNet50                    | $256{	imes}704$   | 2      | 0.372                  | 0.500                  | 0.598                       | 0.270                       | 0.438                       | 0.367                       | 0.190 |
| FB-BEV [14]                 | ResNet50                    | $256{	imes}704$   | 3      | 0.378                  | 0.498                  | 0.620                       | 0.273                       | 0.444                       | 0.374                       | 0.200 |
| SA-BEV [32]                 | ResNet50                    | $256{	imes}704$   | 2      | 0.387                  | 0.512                  | 0.613                       | 0.266                       | 0.352                       | 0.382                       | 0.199 |
| FSD-BEV                     | ResNet50                    | $256{	imes}704$   | 2      | 0.403                  | 0.526                  | 0.576                       | 0.259                       | 0.362                       | 0.359                       | 0.198 |
| $FSD-BEV^{\dagger}$         | ResNet50                    | $256{\times}704$  | 2      | 0.412                  | 0.538                  | 0.527                       | 0.256                       | 0.363                       | 0.330                       | 0.207 |
| DETR3D [26]                 | ResNet101                   | $900{	imes}1600$  | 1      | 0.303                  | 0.374                  | 0.860                       | 0.278                       | 0.437                       | 0.967                       | 0.235 |
| PETR [17]*                  | ResNet101                   | $512{\times}1408$ | 1      | 0.366                  | 0.441                  | 0.717                       | 0.267                       | 0.412                       | 0.834                       | 0.190 |
| BEVFormer [13] <sup>*</sup> | ResNet101                   | $900{	imes}1600$  | 4      | 0.416                  | 0.517                  | 0.673                       | 0.274                       | 0.372                       | 0.394                       | 0.198 |
| BEVDepth [12]               | ResNet101                   | $512{\times}1408$ | 2      | 0.412                  | 0.535                  | 0.565                       | 0.266                       | 0.358                       | 0.331                       | 0.190 |
| Sparse4D $[16]^*$           | ResNet101                   | $900{	imes}1600$  | 4      | 0.436                  | 0.541                  | 0.633                       | 0.279                       | 0.363                       | 0.317                       | 0.177 |
| TiG-BEV [9]                 | ResNet101                   | $512{\times}1408$ | 2      | 0.440                  | 0.544                  | 0.570                       | 0.267                       | 0.392                       | 0.331                       | 0.201 |
| STS [27]                    | ResNet101                   | $512{\times}1408$ | 2      | 0.431                  | 0.542                  | 0.525                       | 0.262                       | 0.380                       | 0.369                       | 0.204 |
| StreamPETR $[23]^*$         | ResNet101                   | $512{\times}1408$ | 8      | 0.504                  | 0.592                  | 0.569                       | 0.262                       | 0.315                       | 0.257                       | 0.199 |
| FSD-BEV                     | $\operatorname{ResNet-101}$ | $512{\times}1408$ | 2      | 0.488                  | 0.589                  | 0.501                       | 0.251                       | 0.316                       | 0.288                       | 0.197 |
| $FSD-BEV^*$                 | $\operatorname{ResNet-101}$ | $512{	imes}1408$  | 2      | 0.500                  | 0.596                  | 0.489                       | 0.251                       | 0.317                       | 0.279                       | 0.204 |

two frames. Further performance gains are observed when we increase the size of distilled BEV features, which exceeds SA-BEV [32] by 2.5% and 2.6% in mAP and NDS. When we use the larger ResNet101 as the backbone, the performance of FSD-BEV is further increased and even exceeds StreamPETR [23], the current state-of-the-art by 0.4% NDS. It is noteworthy that StreamPETR uses much more past frames than FSD-BEV.

The comparison results on the test set are illustrated in Tab. 2. Compared with SOLOFusion [21] that utilizes a larger ConvNeXt-B backbone network and inputs 17 past frames, FSD-BEV achieves a lead of 0.3% and 1.4% in mAP and NDS. When compared to the advanced method SA-BEV [32], which also uses 2 frames, the mAP and NDS of FSD-BEV are improved by 1% and 0.9%.

#### 4.3 Ablation Study

Module Analysis We select BEVDepth as the baseline and conduct training for 24 epochs using the CBGS strategy. We compare the performance of different module combinations. As shown in Tab. 3, when self-distillation is applied individually, significant improvements of 3.1% mAP and 3.0% NDS are achieved. After applying Point Cloud Intensification, there is an obvious improvement in mAP, indicating the enhanced capability of detecting difficult objects. It is consistent with the purpose of this module. When all three modules are integrated, the mAP and NDS are improved by 4% and 4% compared to the baseline. This fully demonstrates the effectiveness of each module.

Table 2: Comparison with the state-of-the-art methods on the nuScenes test set.

| Methods           | Backbone   | Resolution        | Frames | $\mathrm{mAP}\uparrow$ | $\mathrm{NDS}\uparrow$ | $\mathrm{mATE}{\downarrow}$ | $\mathrm{mASE}{\downarrow}$ | mAOE↓ | $\mathrm{mAVE}{\downarrow}$ | $\mathrm{mAAE}{\downarrow}$ |
|-------------------|------------|-------------------|--------|------------------------|------------------------|-----------------------------|-----------------------------|-------|-----------------------------|-----------------------------|
| DETR3D [26]       | V2-99      | $900{	imes}1600$  | 1      | 0.412                  | 0.479                  | 0.641                       | 0.255                       | 0.394 | 0.845                       | 0.133                       |
| UVTR [10]         | V2-99      | $900{	imes}1600$  | 1      | 0.472                  | 0.551                  | 0.577                       | 0.253                       | 0.391 | 0.508                       | 0.123                       |
| BEVFormer [13]    | V2-99      | $900{	imes}1600$  | 4      | 0.481                  | 0.569                  | 0.582                       | 0.256                       | 0.375 | 0.378                       | 0.126                       |
| PETRv2 [18]       | RevCol-L   | $640{\times}1600$ | 2      | 0.512                  | 0.592                  | 0.547                       | 0.242                       | 0.360 | 0.367                       | 0.126                       |
| BEVDepth [12]     | ConvNeXt-B | $640{\times}1600$ | 2      | 0.520                  | 0.609                  | 0.445                       | 0.243                       | 0.352 | 0.347                       | 0.127                       |
| BEVStereo [11]    | V2-99      | $640{\times}1600$ | 2      | 0.525                  | 0.610                  | 0.431                       | 0.246                       | 0.358 | 0.357                       | 0.138                       |
| SA-BEV [32]       | V2-99      | $640{\times}1600$ | 2      | 0.533                  | 0.624                  | 0.430                       | 0.241                       | 0.338 | 0.282                       | 0.139                       |
| SOLOFusion $[21]$ | ConvNeXt-B | $640{\times}1600$ | 17     | 0.540                  | 0.619                  | 0.453                       | 0.257                       | 0.376 | 0.276                       | 0.148                       |
| FSD-BEV           | V2-99      | $640{\times}1600$ | 2      | 0.543                  | 0.633                  | 0.404                       | 0.236                       | 0.325 | 0.286                       | 0.131                       |

Table 3: Performance comparison of different module combinations.

| FSD          | PCI          | MSFE         | $mAP\uparrow$ | $NDS\uparrow$ | mATE↓ | $mASE\downarrow$ | $\mathrm{mAOE}{\downarrow}$ | $\mathrm{mAVE}{\downarrow}$ | $mAAE\downarrow$ |
|--------------|--------------|--------------|---------------|---------------|-------|------------------|-----------------------------|-----------------------------|------------------|
|              |              |              | 0.363         | 0.486         | 0.620 | 0.276            | 0.450                       | 0.387                       | 0.222            |
| $\checkmark$ |              |              | 0.394         | 0.516         | 0.590 | 0.261            | 0.403                       | 0.353                       | 0.200            |
| $\checkmark$ | $\checkmark$ |              | 0.400         | 0.516         | 0.586 | 0.261            | 0.436                       | 0.349                       | 0.202            |
| $\checkmark$ | $\checkmark$ | $\checkmark$ | 0.403         | 0.526         | 0.576 | 0.259            | 0.362                       | 0.359                       | 0.198            |

**Foreground Self-Distillation** We conduct experiments on Foreground Self-Distillation to highlight the benefits of distillation performance brought by foreground segmentation and the results are shown in Tab. 4. It can be found that the student model applying the foreground segmentation outperforms the one without the foreground segmentation by 2.1% mAP, 1.8% NDS. It is partly because only distilling in the foreground region transfers more valuable information, and also because of the giant leap in the teacher's performance (11.6% mAP and 7.7% NDS), telling the importance of a good teacher.

**Point Cloud Intensification** We separately evaluate the accuracy of the teacher branch to assess the impact of different Point Cloud Intensification strategies. As shown in Tab. 5, when employing the FC strategy, we observed improvements of 2.1% in mAP and 0.9% in NDS. However, upon further addition of the PPA strategy, only a slight increase in NDS is observed. This could be attributed to the fact that the FC method might possess similar functionality to PPA, resulting in a reduced gain from PPA.

Multi-Scale Foreground Enhancement Different from pixel-level semantic segmentation tasks, the foreground masks suffice the requirement in FSD-BEV as long as they can roughly indicate the location of foreground objects. A high foreground threshold can filter out more background noise of the scene, but it also leads to possible losses of foreground information. Therefore, we compared the effects of foreground enhancement under different foreground threshold settings and different scale feature maps, as shown in Tab. 6. The results indicate that with a foreground threshold of 0.1, simultaneous use of  $\mathbf{F}^4$  and  $\mathbf{F}^8$  achieves the best enhancement effects.

Table 4: Ablation study of foreground segmentation applied in Foreground Self-Distillation. "FS" denotes the foreground segmentation.

EG D

 
 Table 5: Comparison of the Point Cloud
 Intensification schemes on the teacher branch. "FC" represents Frame Combination, and "PPA" represents Pseudo Point Assignment.

| FS           | Branch  | mAP   | NDS   |
|--------------|---------|-------|-------|
|              | Teacher | 0.468 | 0.564 |
|              | Student | 0.372 | 0.498 |
| $\checkmark$ | Teacher | 0.584 | 0.641 |
| $\checkmark$ | Student | 0.393 | 0.516 |

Table 6: Comparison of multi-scale foreground enhancement effects under different foreground thresholds.  $\beta$  represents the foreground threshold, and  $\mathbf{F}^4$  and  $\mathbf{F}^8$  respectively denote feature maps of different scales (refer to Sec. 3.3).

| $\beta$ | $\mathbf{F}^4$ | $\mathbf{F}^{8}$ | mAP↑  | $NDS\uparrow$ | mATE↓ | $mASE\downarrow$ | $mAOE\downarrow$ | $\mathrm{mAVE}{\downarrow}$ | $mAAE\downarrow$ |
|---------|----------------|------------------|-------|---------------|-------|------------------|------------------|-----------------------------|------------------|
| 0.00    | $\checkmark$   | $\checkmark$     | 0.348 | 0.465         | 0.681 | 0.278            | 0.572            | 0.352                       | 0.214            |
| 0.15    | $\checkmark$   | $\checkmark$     | 0.346 | 0.466         | 0.660 | 0.281            | 0.574            | 0.357                       | 0.202            |
| 0.20    | $\checkmark$   | $\checkmark$     | 0.346 | 0.465         | 0.660 | 0.276            | 0.560            | 0.379                       | 0.209            |
| 0.30    | $\checkmark$   | $\checkmark$     | 0.344 | 0.460         | 0.687 | 0.278            | 0.594            | 0.359                       | 0.207            |
| 0.10    | $\checkmark$   |                  | 0.347 | 0.467         | 0.677 | 0.281            | 0.553            | 0.350                       | 0.207            |
| 0.10    |                | $\checkmark$     | 0.350 | 0.468         | 0.659 | 0.278            | 0.564            | 0.350                       | 0.220            |
| 0.10    | $\checkmark$   | $\checkmark$     | 0.352 | 0.469         | 0.662 | 0.282            | 0.580            | 0.340                       | 0.209            |

#### 4.4 Comparison with Cross-Modal Distillation

We compare FSD-BEV with methods that apply cross-modal distillation to improve the camera-only BEV-based 3D detectors. From the results in Tab. 7, it can be found that the simple foreground self-distillation framework outperforms other distillation methods that adopt complicated strategies. The advantage of FSD-BEV is more significant when the ResNet101 is chosen as the image backbone and a total of 7.6% and 5.4% in mAP and NDS is achieved, which largely exceeds the distillation gains of other methods.

The joint progress of the student model and the teacher model is a benefit characteristic of FSD-BEV. To evaluate the profit of this characteristic, we chose the state-of-the-art cross-modal distillation approach, DistillBEV [28], as the control group. Each model is configured with two sets of control experiments, the model in one set using ResNet50 as the backbone network and the models in the other set using ResNet101. The comparison results are shown in Tab. 8.

Under the setting of ResNet50, the student accuracy of FSD-BEV is roughly comparable to that of DistillBEV. However, when we switch to a larger backbone, the student accuracy of DistillBEV improves by 4.7% and 3.7% in mAP and NDS, FSD-BEV improves by 8.5% and 6.3%. It is worth noting that our teacher does not exhibit higher performance compared to the teacher of Distill-BEV, indicating our method's higher distillation efficiency. On the other hand, compared to cross-modal methods with a fixed teacher accuracy, our teacher accuracy increased by 9.4% in mAP and 7.6% in NDS when a larger backbone network is employed. This growth aligns with the magnitude of improvement

**Table 7:** Comparison with cross-modal distillation methods. The improvements caused by distillation are shown in red.

| Methods    | Backbone  | Teacher Modality | Student  | mAP                   | NDS                         |
|------------|-----------|------------------|----------|-----------------------|-----------------------------|
| BEVDistill | ResNet50  | LiDAR            | BEVDepth | 0.330(+1.3%)          | 0.452(+1.2%)                |
| UniDistill | ResNet50  | LiDAR+Camera     | BEVDet   | 0.296(+3.2%)          | 0.393(+ <mark>3.2%</mark> ) |
| FSD-BEV    | ResNet50  | LiDAR+Camera     | BEVDepth | 0.403( <b>+4.0%</b> ) | 0.526(+4.0%)                |
| TiG-BEV    | ResNet101 | LiDAR            | BEVDepth | 0.440(+2.4%)          | 0.544(+2.3%)                |
| $X^{3}KD$  | ResNet101 | LiDAR            | BEVDepth | 0.448(+ <b>3.9%</b> ) | 0.553(+2.2%)                |
| DistillBEV | ResNet101 | LiDAR+Camera     | BEVDepth | 0.450(+3.8%)          | 0.547(+1.2%)                |
| FSD-BEV    | ResNet101 | LiDAR+Camera     | BEVDepth | 0.488(+7.6%)          | 0.589(+5.4%)                |

Table 8: Comparison of distillation performance for different model magnitudes.

| Methods         | Backbone  | Tea   | cher  | Student          |                     |  |
|-----------------|-----------|-------|-------|------------------|---------------------|--|
| wicthous        | Dackbolic | mAP   | NDS   | mAP              | NDS                 |  |
| DistillBEV [28] | ResNet50  | 0.671 | 0.780 | 0.403            | 0.510               |  |
|                 | ResNet101 | 0.671 | 0.780 | $0.450_{+4.7\%}$ | $0.547_{+3.7\%}$    |  |
| Ours            | ResNet50  | 0.584 | 0.641 | 0.403            | 0.526               |  |
| Ours            | ResNet101 | 0.680 | 0.717 | $0.488_{+8.5\%}$ | $0.589_{\pm 6.3\%}$ |  |

seen in the student. Such cooperative growth enables our approach to maintain stable distillation performance.

# 5 Conclusion

In order to narrow the performance gap between LiDAR-based 3D detectors and BEV-based multi-view 3D detectors, several methods utilize cross-modal knowledge distillation attempts to transfer beneficial information from the extra teacher model to the student model. However, such approaches always suffer from feature distribution discrepancies originating from different data modalities and network structures. In this paper, we propose a Foreground Self-Distillation (FSD) scheme that effectively alleviates this issue without the need for pretrained teacher models or cumbersome distillation strategies. Additionally, we design two Point Cloud Intensification (PCI) strategies to compensate for the sparsity of point clouds by frame combination and pseudo point assignment, further intensifying the performance of the teacher model. Finally, we develop a Multi-Scale Foreground Enhancement (MSFE) module to extract and fuse multi-scale foreground features by predicted elliptical Gaussian heatmap, further improving the model's performance.

# Acknowledgement

This work was supported by Zhejiang Provincial Natural Science Foundation of China (No. LD24F020016), "Pioneer" and "Leading Goose" R&D Program of Zhejiang (No. 2024C01020), the National Natural Science Foundation of China (No. 62176017, No. 62302031), and ZEEKR Intelligent Technology Co., Ltd.

15

# References

- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020) 9, 10
- Chen, Z., Li, Z., Zhang, S., Fang, L., Jiang, Q., Zhao, F.: Bevdistill: Cross-modal bev distillation for multi-view 3d object detection. arXiv preprint arXiv:2211.09386 (2022) 2, 4, 5
- Dai, X., Jiang, Z., Wu, Z., Bao, Y., Wang, Z., Liu, S., Zhou, E.: General instance distillation for object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7842–7851 (2021) 4
- Guo, J., Han, K., Wang, Y., Wu, H., Chen, X., Xu, C., Xu, C.: Distilling object detectors via decoupled features. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2154–2164 (2021) 4
- Guo, X., Shi, S., Wang, X., Li, H.: Liga-stereo: Learning lidar geometry aware representations for stereo-based 3d detector. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3153–3163 (2021) 4
- Hou, Y., Ma, Z., Liu, C., Hui, T.W., Loy, C.C.: Inter-region affinity distillation for road marking segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12486–12495 (2020) 4
- Huang, J., Huang, G.: Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. arXiv preprint arXiv:2203.17054 (2022) 11
- Huang, J., Huang, G., Zhu, Z., Ye, Y., Du, D.: Bevdet: High-performance multicamera 3d object detection in bird-eye-view. arXiv preprint arXiv:2112.11790 (2021) 4, 11
- Huang, P., Liu, L., Zhang, R., Zhang, S., Xu, X., Wang, B., Liu, G.: Tig-bev: Multiview bev 3d object detection via target inner-geometry learning. arXiv preprint arXiv:2212.13979 (2022) 2, 5, 11
- Li, Y., Chen, Y., Qi, X., Li, Z., Sun, J., Jia, J.: Unifying voxel-based representation with transformer for 3d object detection. Advances in Neural Information Processing Systems 35, 18442–18455 (2022) 4, 12
- Li, Y., Bao, H., Ge, Z., Yang, J., Sun, J., Li, Z.: Bevstereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1486–1494 (2023) 4, 11, 12
- Li, Y., Ge, Z., Yu, G., Yang, J., Wang, Z., Shi, Y., Sun, J., Li, Z.: Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1477–1485 (2023) 4, 6, 11, 12
- Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Qiao, Y., Dai, J.: Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In: European conference on computer vision. pp. 1–18. Springer (2022) 4, 11, 12
- Li, Z., Yu, Z., Wang, W., Anandkumar, A., Lu, T., Alvarez, J.M.: Fb-bev: Bev representation from forward-backward view transformations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6919–6928 (2023) 11
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2117–2125 (2017) 9

- 16 Z. Jiang, J. Zhang et al.
- Lin, X., Lin, T., Pei, Z., Huang, L., Su, Z.: Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion. arXiv preprint arXiv:2211.10581 (2022) 11
- Liu, Y., Wang, T., Zhang, X., Sun, J.: Petr: Position embedding transformation for multi-view 3d object detection. In: European Conference on Computer Vision. pp. 531–548. Springer (2022) 4, 11
- Liu, Y., Yan, J., Jia, F., Li, S., Gao, A., Wang, T., Zhang, X.: Petrv2: A unified framework for 3d perception from multi-camera images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3262–3272 (2023) 4, 11, 12
- 19. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017) 10
- Park, D., Ambrus, R., Guizilini, V., Li, J., Gaidon, A.: Is pseudo-lidar needed for monocular 3d object detection? In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3142–3152 (2021) 2
- Park, J., Xu, C., Yang, S., Keutzer, K., Kitani, K., Tomizuka, M., Zhan, W.: Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. arXiv preprint arXiv:2210.02443 (2022) 11, 12
- Philion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16. pp. 194–210. Springer (2020) 4
- Wang, S., Liu, Y., Wang, T., Li, Y., Zhang, X.: Exploring object-centric temporal modeling for efficient multi-view 3d object detection. arXiv preprint arXiv:2303.11926 (2023) 4, 11
- Wang, T., Xinge, Z., Pang, J., Lin, D.: Probabilistic and geometric depth: Detecting objects in perspective. In: Conference on Robot Learning. pp. 1475–1485. PMLR (2022) 2
- Wang, T., Zhu, X., Pang, J., Lin, D.: Fcos3d: Fully convolutional one-stage monocular 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 913–922 (2021)
- Wang, Y., Guizilini, V.C., Zhang, T., Wang, Y., Zhao, H., Solomon, J.: Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In: Conference on Robot Learning. pp. 180–191. PMLR (2022) 4, 11, 12
- Wang, Z., Min, C., Ge, Z., Li, Y., Li, Z., Yang, H., Huang, D.: Sts: Surroundview temporal stereo for multi-view 3d detection. arXiv preprint arXiv:2208.10145 (2022) 4, 11
- Wang, Z., Li, D., Luo, C., Xie, C., Yang, X.: Distillbev: Boosting multi-camera 3d object detection with cross-modal knowledge distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8637–8646 (2023) 2, 4, 5, 13, 14
- Yang, C., Xie, L., Su, C., Yuille, A.L.: Snapshot distillation: Teacher-student optimization in one generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2859–2868 (2019) 4
- 30. Yang, C., Chen, Y., Tian, H., Tao, C., Zhu, X., Zhang, Z., Huang, G., Li, H., Qiao, Y., Lu, L., et al.: Bevformer v2: Adapting modern image backbones to bird's-eye-view recognition via perspective supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17830–17839 (2023) 4
- Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11784–11793 (2021) 1

- Zhang, J., Zhang, Y., Liu, Q., Wang, Y.: Sa-bev: Generating semantic-aware bird's-eye-view feature for multi-view 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3348–3357 (2023) 3, 6, 7, 11, 12
- 33. Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., Ma, K.: Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3713–3722 (2019) 4
- 34. Zhang, Y., Huang, D., Wang, Y.: Pc-rgnn: Point cloud completion and graph neural network for 3d object detection. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 3430–3437 (2021) 1
- Zhou, C., Zhang, Y., Chen, J., Huang, D.: Octr: Octree-based transformer for 3d object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5166–5175 (2023) 1
- Zhou, S., Liu, W., Hu, C., Zhou, S., Ma, C.: Unidistill: A universal cross-modality knowledge distillation framework for 3d object detection in bird's-eye view. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5116–5125 (2023) 2, 4, 5
- 37. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019) 9
- Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4490–4499 (2018) 1
- Zhu, B., Jiang, Z., Zhou, X., Li, Z., Yu, G.: Class-balanced grouping and sampling for point cloud 3d object detection. arXiv preprint arXiv:1908.09492 (2019) 10