

DistillBEV: Boosting Multi-Camera 3D Object Detection with Cross-Modal Knowledge Distillation

Zeyu Wang^{1,2*} Dingwen Li^{1*} Chenxu Luo¹ Cihang Xie² Xiaodong Yang^{1†}
¹QCraft ²UC Santa Cruz

Abstract

3D perception based on the representations learned from multi-camera bird’s-eye-view (BEV) is trending as cameras are cost-effective for mass production in autonomous driving industry. However, there exists a distinct performance gap between multi-camera BEV and LiDAR based 3D object detection. One key reason is that LiDAR captures accurate depth and other geometry measurements, while it is notoriously challenging to infer such 3D information from merely image input. In this work, we propose to boost the representation learning of a multi-camera BEV based student detector by training it to imitate the features of a well-trained LiDAR based teacher detector. We propose effective balancing strategy to enforce the student to focus on learning the crucial features from the teacher, and generalize knowledge transfer to multi-scale layers with temporal fusion. We conduct extensive evaluations on multiple representative models of multi-camera BEV. Experiments reveal that our approach renders significant improvement over the student models, leading to the state-of-the-art performance on the popular benchmark nuScenes.

1. Introduction

Perceiving 3D environments is essential for autonomous driving as it is crucial for subsequent onboard modules from prediction [28, 37] to planning [2, 18]. Although LiDAR based methods have achieved remarkable progress [15, 25, 26, 42], there have recently been fast growing attentions to camera based approaches from both academia and industry [1, 32]. Compared to a sensor suite of LiDAR, the cost of cameras is typically 10 times cheaper, gaining a huge cost advantage for mass-production OEMs. Additionally, cameras are better suited to detect distant objects and recognize visual based road elements (e.g., traffic lights and signs).

A straightforward solution for camera based 3D object detection is the monocular paradigm that has been broadly




Figure 1. An overview of performance improvement in terms of mAP and NDS on the validation set of nuScenes. Enabled by the proposed cross-modal knowledge distillation method DistillBEV, a variety of multi-camera BEV based 3D object detectors achieve consistent and significant performance boost.

studied [27, 30, 36]. However, such methods have to deal with surrounding cameras separately and require complex post-processing to fuse detections from multiple views. As an alternative, the bird’s-eye-view (BEV) based framework is drawing extensive attentions to offer a holistic feature representation space from multi-camera images, and has made substantial improvement [11, 12, 20, 21]. BEV based framework owns the following inherent merits including (i) joint feature learning from multi-view images, (ii) unified detection space without post fusion, (iii) amenability for temporal fusion, and (iv) convenient output representation for the downstream prediction and planning.

Despite the advancement achieved in this field, a distinct performance gap remains between multi-camera BEV and LiDAR based 3D object detection. For instance, the leading multi-camera BEV based method is outperformed by its LiDAR counterpart over 15% mAP and 10% NDS on popular benchmark nuScenes [3]. On the other hand, a data collection fleet can be equipped with both cameras and LiDAR, while the mass-produced vehicles can be LiDAR-free.

In light of above observations, we present **DistillBEV**: a simple and effective cross-model knowledge distillation approach to bridge the feature learning gap between multi-camera BEV and LiDAR based detectors. Our strategy to

*Equal contribution

†Correspondence to xiaodong@qcrafft.ai

ward achieving this goal is to align the corresponding features learned from images and point clouds. Given merely images as input, the multi-camera BEV based detector (student) is guided to imitate the features extracted from point clouds by a well-trained LiDAR based detector (teacher). We argue that the accurate 3D geometric cues such as depth and shape as well as how such cues are represented in the point cloud features provide valuable guidance to the training process of student model. Moreover, we emphasize that our approach incurs no extra computation cost during inference as the teacher model involves in training only.

Due to the notable discrepancy between the two modalities, the cross-modal knowledge distillation is extremely challenging. Compared to camera images, LiDAR scans are inherently sparse, and the majority of 3D space is empty. Even if in the occupied space, background (e.g., building and vegetation) dominates, and objects of interest (e.g., bus and pedestrian) come of varied sizes. It is thus nontrivial to locate the informative regions to make the knowledge transfer more focused, and meanwhile balance the distillation importance assigned to different foreground objects. Furthermore, the teacher and student models are respectively developed in their specific domains, leading to disparate network architectures. It is also demanding to generalize the cross-modal distillation to adapt to various combinations of different teacher and student networks.

To tackle these challenges for DistillBEV, we first propose region decomposition to partition a feature map into true positive, false positive, true negative and false negative regions, which elaborately decouple foreground and background in the cross-modal distillation. Based on the decomposition, we introduce adaptive scaling to balance the significantly varied box sizes when objects are presented in BEV. We further exploit spatial attention to encourage the student model to mimic the attention pattern produced by the teacher model, so as to focus on the crucial features for effective knowledge transfer. We then extend the distillation to multi-scale layers to achieve thorough feature alignment between teacher and student. Finally, we incorporate temporal information for both teacher and student models in BEV, thus enabling the distillation with temporal fusion readily. Thanks to the proposed generalizable design, our approach is flexible to be applied to various combinations of teacher and student detectors. As illustrated in Figure 1, DistillBEV consistently and remarkably improves multiple representative student models.

Our main contributions are summarized as follows. First, we present the cross-modal distillation in BEV, which naturally suits for knowledge transfer between LiDAR and multi-camera BEV based detectors. Second, we propose the effective balancing design to enable the student to focus on learning crucial features of the teacher with multiple scales and temporal fusion. Third, our approach achieves superior

performance, and more importantly, obtains consistent and considerable improvement on various teacher-student combinations. Our code and model will be made available at <https://github.com/qcraftai/distill-bev>.

2. Related Work

Camera based 3D Object Detection. A large family of the methods in this field are based on the monocular paradigm, such as FCOS3D [36] and DD3D [30], which resemble 2D object detection. Recently, the multi-camera BEV based framework is trending thanks to its inherent merits. In this framework, the view transformation module plays a fundamental role to convert multi-view image features to BEV. Some methods employ inverse perspective mapping [4] or multilayer perceptron [29] to perform the translation from perspective view to BEV. In [32] LSS is introduced to lift image features by corresponding bin-wise depth distribution. This line of research includes a series of works such as BEVDet [12], BEVDet4D [11], and BEVDepth [20]. BEVFormer is proposed in [21] to utilize cross-attention to look up and aggregate image features across cameras. In addition, BEV representations provide a more desirable connection of scene features at multiple timestamps to improve object detection and motion state estimation. BEVDet4D and BEVDepth fuse previous and current features using spatial alignment, and BEVFormer performs fusion through temporal attention in a soft way.

LiDAR based 3D Object Detection. Since most methods of this field apply voxelization to transform irregular point clouds to regular grids such as pillars or voxels, it is natural to extract features in BEV. VoxelNet [45] is a pioneering approach using 3D convolutions on the voxel features extracted by [33]. SECOND [40] introduces sparse 3D convolutions to improve computation efficiency. PointPillars [15] proposes to collapse height dimension and use 2D convolutions to further reduce inference latency. CenterPoint [42] is a popular anchor-free method that represents objects as points. PillarNeXt [17] shows that the pillar based models with modernized designs in architecture and training outperforms the voxel counterpart in both accuracy and latency. Fusing multiple sensors is also widely used to enhance the detection performance. MVP [43] is a sensor-fusion version of CenterPoint enhanced by image virtual points.

Knowledge Distillation. This technique is originally proposed for network compression by transferring the information from a large teacher model to a compact student model [9]. Most methods are initially designed for image classification, but hardly make improvement for image object detection. Some recent methods have successfully adapted knowledge distillation for 2D object detection [7, 13, 41, 44]. Nevertheless, the research on distillation for 3D object detection has been less explored, in particular when the teacher and student models come from dif-

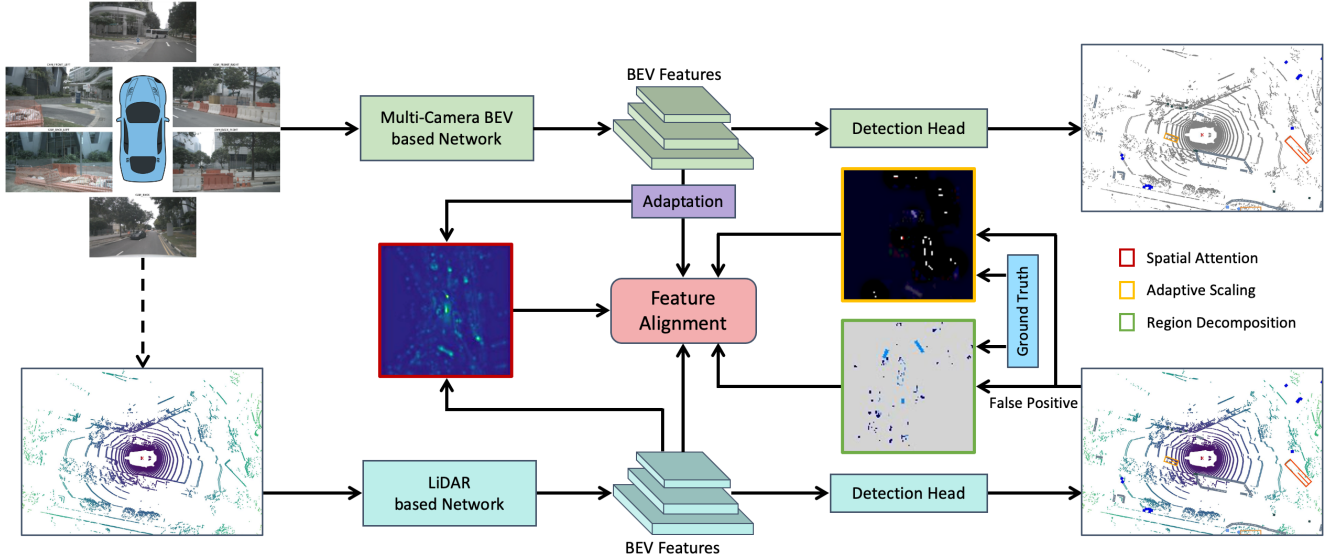


Figure 2. A schematic overview of the proposed cross-modal distillation approach DistillBEV. We aim to bridge the representation learning gap between multi-camera BEV (student) and LiDAR (teacher) based detectors by guiding the former to imitate the features extracted by the latter. We introduce the balancing strategy including region decomposition, adaptive scaling and spatial attention to encourage the student to focus on the crucial features to align with the teacher. Our teacher model can be based on either LiDAR or camera-LiDAR fusion (indicated by the dashed line). Note the bottom branch of LiDAR related components are removed after training.

ferent modalities [6, 19]. The most related work to our proposed approach is [5], which introduces a dense foreground-guided feature imitation and a sparse instance-wise distillation to transfer spatial knowledge from LiDAR to multi-camera 3D object detection. Compared to this method, our approach achieves more fine-grained distillation through introducing region decomposition and adaptive scaling. Additionally, our design accommodates multi-scale distillation, which enhances cross-modal knowledge transfer at different feature abstractions. Experimental results also confirm the superiority of our approach.

3. Method

As illustrated in Figure 2, to tackle the challenges for cross-modal knowledge transfer, DistillBEV involves formulating region decomposition mask, adaptive scaling factor and spatial attention map, as well as making extensions to multi-scale layers and temporal fusion.

3.1. Region Decomposition

It has been well acknowledged in 2D object detection that simply performing feature alignment between teacher and student hardly makes improvement due to the imbalance between foreground and background regions [7, 13, 41, 44]. However, this phenomenon is even more severe in 3D object detection as the vast majority of 3D space is unoccupied. Our statistics on BEV feature maps finds that less than 30% pixels are non-empty on average, and

among them, only a small fraction contains objects that we are interested in. To perform effective knowledge transfer, we introduce the region decomposition to guide the student to focus on crucial regions rather than treating all regions equally. Specifically, we partition a feature map into four types: true positive (TP), false positive (FP), true negative (TN) and false negative (FN). Accordingly, we define a region decomposition mask M :

$$M_{i,j} = \begin{cases} 1, & \text{if } (i,j) \in \text{TP or FN} \\ \eta, & \text{if } (i,j) \in \text{FP} \\ 0, & \text{if } (i,j) \in \text{TN} \end{cases} \quad (1)$$

where i, j are the coordinates on a feature map, and η controls the relative importance of pixels in FP regions.

This decomposition gives our approach the flexibility in assigning varied importance on different regions. It is straightforward to take into account the regions covered by ground truth boxes (i.e., the union of TP and FN regions), which exactly convey the features from foreground objects. However, we also treat FP regions differently from TN regions. When the teacher model generates high activations on some regions even though they are FP (e.g., a pole is mis-detected as a pedestrian), encouraging the student model to mimic such feature responses can be still beneficial for the overall 3D geometry learning. FP regions can be found by thresholding the confidence heatmaps produced by the teacher detector and ground truth labels:

$$\text{FP} = (H^t > \gamma) \& (H^g < \gamma), \quad (2)$$

where H^t and H^g correspond to the heatmaps obtained from the teacher model and ground truth, respectively, and γ is a hyper-parameter for heatmap thresholding.

3.2. Adaptive Scaling

Another challenge to distill knowledge from teacher to student in BEV is the great span of various object sizes. For instance, a bus is dozens of times the size of a pedestrian from bird’s-eye-view. Moreover, background such as walls and plants overwhelm the non-empty regions. Thus, background stuff and giant foreground objects would dominate the distillation loss since substantially more features come from them. It is desirable to reflect different-sized objects or classes with similar contributions in the distillation loss. An adaptive scaling factor is introduced to achieve this goal:

$$S_{i,j} = \begin{cases} \frac{1}{\sqrt{H_k W_k}}, & \text{if } (i,j) \in O_k \\ \frac{1}{N_{\text{FP}}}, & \text{if } (i,j) \in \text{FP} \\ \frac{1}{N_{\text{TN}}}, & \text{if } (i,j) \in \text{TN} \end{cases} \quad (3)$$

where O_k is the k -th ground truth object (TP or FN) with bounding box length H_k and width W_k in BEV, N_{FP} and N_{TN} denote the number of pixels falling in the FP and TN regions, respectively.

3.3. Spatial Attention

Attention maps [10, 41, 44] have been applied to various tasks and architectures to improve representation learning by focusing on the important features as well as suppressing the insignificant ones. Here we adopt a spatial attention map based on the extracted features of both teacher and student to further select more informative features to concentrate on. A spatial attention map is constructed by:

$$\mathcal{P}(F)_{i,j} = \frac{1}{C} \sum_{c=1}^C |F_{c,i,j}|, \quad (4)$$

$$\mathcal{N}(F) = HW \text{ softmax}(\mathcal{P}(F)/\tau), \quad (5)$$

where $F \in \mathbb{R}^{C \times H \times W}$ is a feature map, $\mathcal{P}(F) \in \mathbb{R}^{H \times W}$ denotes the average pooling result of the absolute values of F along the channel dimension, $\mathcal{N}(F) \in \mathbb{R}^{H \times W}$ is the normalized attention by softmax over all spatial locations, and τ is a temperature to adjust the distribution entropy. We obtain the final spatial attention map by considering the feature maps from both teacher F^t and student F^s :

$$A = (\mathcal{N}(F^t) + \mathcal{N}(\tilde{F}^s))/2, \quad \tilde{F}^s = \mathcal{G}(F^s), \quad (6)$$

where \mathcal{G} is an adaptation module to map F^s to \tilde{F}^s that is with the same size as F^t . More details of the design choices about \mathcal{G} are discussed in Sections 3.4 and 4.3.

3.4. Multi-Scale Distillation

It is commonly believed that layers of different depth in a network encodes varying feature abstractions [14]. One successful application is the feature pyramid network [22], which combines features from different levels to better detect objects of various sizes [17, 35, 42]. To realize comprehensive alignment between teacher and student, we adopt this idea to perform feature distillation at multiple scales for the models based on CNNs. However, the teacher and student networks are separately designed with different architectures, making it nontrivial to find intermediate feature correspondence. For instance, the BEV feature maps in teacher are usually $2 \times$ or $4 \times$ the size of those in student. Naively aligning features of the same resolution results in incompatibility of the feature abstraction level. Hence, we introduce a lightweight adaptation module \mathcal{G} consisting of upsampling and projection layers to map the student feature before aligning with the teacher feature at the similar level. We also find that feature imitation at early layers is detrimental to the distillation, which is because the representation discrepancy caused by modality gap between point clouds and images remains substantial at the early stage. Note we identify and utilize the FP regions only at the last encoding layer of BEV (i.e., the pre-head feature). We find this setting works the best presumably because FP regions are better expressed by high-level semantic features emerging at the last layer. See more architecture details in the supplementary material.

3.5. Distillation Loss

We train a student network with the original loss including classification and regression as well as the overall distillation loss, the related terms of which are summarized as follows. We first define the feature imitation loss at the n -th distillation layer between teacher $F^{t(n)}$ and student $F^{s(n)}$:

$$L_{\text{feat}} = \alpha \sum_{c=1}^C \sum_{i=1}^H \sum_{j=1}^W M_{i,j} S_{i,j} A_{i,j} \left(F_{c,i,j}^{t(n)} - \tilde{F}_{c,i,j}^{s(n)} \right)^2 + \beta \sum_{c=1}^C \sum_{i=1}^H \sum_{j=1}^W \bar{M}_{i,j} S_{i,j} A_{i,j} \left(F_{c,i,j}^{t(n)} - \tilde{F}_{c,i,j}^{s(n)} \right)^2, \quad (7)$$

where $\tilde{F}^{s(n)} = \mathcal{G}(F^{s(n)})$, \bar{M} is the logical complement of the region decomposition mask M , S denotes the adaptive scaling factor, A is the spatial attention map, and α and β are the hyper-parameters to weight the two terms.

Furthermore, we exploit an attention imitation loss to enforce the student to learn to generate an attention pattern similar to the teacher, and therefore focus on the spatial locations that the teacher network considers more crucial:

$$L_{\text{attn}} = \sum_{i=1}^H \sum_{j=1}^W \left| \mathcal{P}(F^{t(n)})_{i,j} - \mathcal{P}(\tilde{F}^{s(n)})_{i,j} \right|. \quad (8)$$

In summary, the overall distillation objective function is the sum of feature imitation loss (7) and attention imitation loss (8) at multiple scales:

$$L_{\text{dist}} = \sum_{n=1}^N L_{\text{feat}}(F^{t(n)}, F^{s(n)}) + \lambda L_{\text{attn}}(F^{t(n)}, F^{s(n)}), \quad (9)$$

where N is the number of selected layers to perform distillation, and λ controls the relative importance between the two loss functions.

3.6. Distillation with Temporal Fusion

One desirable property of the representations in multi-camera BEV is to facilitate fusing of features from multiple timestamps. Methods [11, 20, 21] developed with temporal fusion greatly improve 3D object detection and motion estimation by leveraging important dynamic cues. As for LiDAR based models, it is a common practice to fuse multiple point clouds by directly transforming the past sweeps to the current coordinate frame through ego-motion compensation, and a relative timestamp is added to the measurements of each point [42]. Thus, it is natural to conduct temporal knowledge transfer in our approach as the teacher can be readily compatible with the student making use of temporal information. In practice, we employ a unified teacher model for both single-frame and multi-frame based student models to perform distillation with temporal fusion.

4. Experiments

In this section, we describe the evaluation setup including dataset, metrics, and implementation details. A variety of ablation study and related analysis are conducted for in-depth understanding of each individual component in our approach. We report extensive comparisons with the state-of-the-art methods on the popular benchmark.

4.1. Experimental Setup

Dataset and Metrics. We evaluate our approach on the large-scale autonomous driving benchmark nuScenes [3]. This dataset consists of 1,000 scenes of roughly 20 seconds each, captured by a 32-beam LiDAR and 6 cameras at the frequency of 20Hz and 10Hz. There are 10 classes in total for 3D object detection, and the annotations are provided at 2Hz. Following the standard evaluation split, 700, 150 and 150 scenes are respectively used for training, validation and test. We follow the official evaluation metrics including mean average precision (mAP) and nuScenes detection score (NDS) as the main metrics. We also use mATE, mASE, mAOE, mAVE and mAAE to measure translation, scale, orientation, velocity and attribute related errors.

Teacher and Student Models. To validate the generalizability of our approach, we consider various teacher and student models. We employ the popular CenterPoint [42]

or its sensor-fusion version MVP [43] as the teacher model. More details are in the supplementary material. As for the student model, we choose BEVDet [12], BEVDet4D [11], BEVDepth [20], and BEVFormer [21] as the representative student models, which represent a broad range of student models from CNNs to Transformers, as well as from the basic version to the temporal (“4D” to incorporate temporal fusion) and spatial (“Depth” to enhance trustworthy depth estimation) extensions. These models together form 8 different teacher-student combinations.

Implementation Details. We implement our approach in PyTorch [31], and train the networks by using 8 NVIDIA Tesla V100 GPUs with the batch size of 64. AdamW [24] is adopted as the optimizer with a cosine-scheduled learning rate of $2e-4$. All models are trained for 24 epochs with the strategy of CBGS [46]. Following [12, 20], data augmentations are applied in both image and BEV spaces. We follow the standard evaluation protocol to set the detection range to $[-51.2\text{m}, 51.2\text{m}] \times [-51.2\text{m}, 51.2\text{m}]$. ResNet-50 [8] pre-trained on ImageNet-1K is used as image backbone and image size is processed to 256×704 , unless otherwise specified. We adopt the common inheriting practice [13] to initialize the detection head of student by the parameters of teacher for faster convergence. More details can be found in the supplementary material.

4.2. Comparison with State-of-the-Art Methods

We start from comparing the student performance before and after distillation on the validation set of nuScenes. As shown in Table 1 and Figure 1, DistillBEV using different teacher models considerably and consistently boosts the four representative students over various metrics. In particular, the most significant performance gains are obtained on BEVFormer, i.e., 4.4% mAP and 4.2% NDS. For the leading algorithm BEVDepth, we also observe the improvement of 3.9% mAP and 2.6% NDS, revealing that our distillation effect is not diminishing with the stronger student.

Taking a closer look into other metrics, we observe that DistillBEV largely improves mAVE of the single-frame based model (i.e., BEVDet). We attribute this to the temporal fusion nature in the teacher models, which effectively transfer temporal knowledge to the distilled students via our approach. And the trend of mAAE follows mAVE as this attribute is predicted based on the velocity estimation. In general, there exists clear improvement in mATE and mAOE, thanks to the accurate depth and geometric cues encoded in point clouds. It is found that mASE is only slightly improved as the students already estimate object scale reasonably well by using visual information alone.

We next compare DistillBEV with the state-of-the-art methods on the validation set and test set of nuScenes. Table 2 shows that, on the validation set, our approach outperforms the camera based methods (2nd group) by a clear

Teacher	Mode	Student	Mode	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
-	-	BEVDet	C	30.5	37.8	72.1	27.9	57.9	91.4	25.0
CenterPoint	L	BEVDet	C	32.7	40.7	70.9	26.5	56.5	81.2	21.0
MVP	L&C	BEVDet	C	34.0	41.6	70.4	26.6	55.6	81.5	20.1
-	-	BEVDet4D	C	32.8	45.9	69.5	27.9	50.8	36.5	20.6
CenterPoint	L	BEVDet4D	C	36.3	48.4	66.6	26.8	49.8	34.9	19.9
MVP	L&C	BEVDet4D	C	37.0	48.8	67.6	26.8	46.1	36.8	20.0
-	-	BEVDepth	C	36.4	48.4	64.9	27.3	49.8	34.9	20.7
CenterPoint	L	BEVDepth	C	38.9	49.8	63.0	26.7	50.4	36.0	20.2
MVP	L&C	BEVDepth	C	40.3	51.0	62.3	26.6	46.4	35.7	20.7
-	-	BEVFormer	C	32.3	43.4	79.6	28.3	53.1	46.0	21.4
CenterPoint	L	BEVFormer	C	35.5	46.8	71.9	27.7	50.8	39.3	20.0
MVP	L&C	BEVFormer	C	36.7	47.6	72.1	27.5	50.6	37.6	20.0

Table 1. Comparison of our approach using various combinations of teacher and student models on the validation set of nuScenes. ‘‘C’’ and ‘‘L’’ indicate the modality of camera and LiDAR, respectively.

Method	Backbone	Mode	mAP	NDS
CenterPoint [42]	-	L	56.4	64.8
MVP [43]	-	L&C	67.1	78.0
FCOS3D [36]	R101	C	34.3	41.5
PETR [23]	R101	C	35.7	42.1
DETR3D [38]	R101	C	34.6	42.5
BEVFormer [21]	R50	C	32.3	43.4
BEVFormer [21]	R101	C	41.6	51.7
BEVDepth [20]	R50	C	35.1	47.5
BEVDepth [20]	R101	C	41.2	53.5
Set2Set [39]	R50	C	37.5	47.9
FitNet [34]	R50	C	37.3	48.0
MonoDistill [6]	R50	C	39.0	49.5
UVTR [19]	R50	C	39.4	50.1
BEVDistill [5]	R50	C	40.7	51.5
BEVDistill [5]	R101	C	41.7	52.4
Ours (BEVFormer)	R50	C	36.7	47.6
Ours (BEVFormer)	R101	C	44.6	54.5
Ours (BEVDepth)	R50	C	40.3	51.0
Ours (BEVDepth)	R101	C	45.0	54.7

Table 2. Comparison on the validation set of nuScenes. Groups 1-4 correspond to the teacher models, the camera based works, the distillation based methods, and our proposed approach.

margin with the same backbone settings. Comparing with other knowledge distillation based methods (3rd group), our approach also performs better than various competing algorithms. As for the test set, we follow [20] to increase the input image size to 640×1600 and double the grid size to 256×256 in BEV, and the pre-training on ImageNet-22K is further applied. As shown in Table 3, our approach achieves superior performance without bells and whistles, such as model ensembling and test-time augmentation.

4.3. Ablation Study

We perform extensive experiments to analyze and understand the proposed components and related design choices in DistillBEV. Unless otherwise mentioned, we disable the

Method	Backbone	Mode	mAP	NDS
CenterPoint [42]	-	L	58.0	65.5
MVP [43]	-	L&C	66.4	70.5
FCOS3D [36]	R101	C	35.8	42.8
BEVDet [12]	Swin-B	C	39.8	46.3
DD3D [30]	VoV-99	C	41.8	47.7
DETR3D [38]	VoV-99	C	41.2	47.9
PETR [23]	VoV-99	C	44.1	50.4
BEVDet4D [12]	Swin-B	C	45.1	56.9
BEVFormer [21]	VoV-99	C	48.1	56.9
BEVDistill [5]	ConvNeXt-B	C	49.8	59.4
BEVDepth [20]	VoV-99	C	50.3	60.0
Ours (BEVDepth)	Swin-B	C	52.5	61.2

Table 3. Comparison on the test set of nuScenes. Groups 1-3 are the teacher models, the state-of-the-art methods, and our proposed approach. Note: VoV-99 [16] is pre-trained on depth prediction with extra data of 15M images and paired point clouds [30].

head inheriting, and exemplify the teacher and student detectors using MVP and BEVDet4D, respectively.

Contribution of Individual Component. Here we study the contribution of each individual component in our design through a set of combination experiments. Starting from the basic feature imitation at multiple scales, we gradually add foreground and background region decomposition (no FP), spatial attention, adaptive scaling, and FP region decomposition. As shown in Table 4d, the basic feature imitation barely improves the baseline (32.8% vs 32.9% mAP, 45.9% vs 45.6% NDS), indicating that the naive feature alignment is inadequate for cross-modal distillation in BEV. In the proposed approach, each component makes decent improvement upon the basic feature imitation, and in the full model, these components collectively boost the baseline result by 5.0% mAP and 3.1% NDS.

Where to Distill. Next we study the impact of distilling at different layers between teacher and student in DistillBEV. As mentioned in Section 3.4, FP regions are only used on the pre-head feature map. So we disable the FP region de-

Model	Threshold	mAP	NDS
Teacher	0.05	36.9	48.4
	0.1	37.8	49.0
	0.3	37.6	49.0
Student	0.1	37.5	48.7

(a) Evaluation of different FP region mining principles by teacher and student under various thresholds.

FP	mAP	NDS
Pre-Head (H)	37.8	49.0
All (H-B2-B1)	37.4	48.6

(b) Evaluation of the impacts of FP region decomposition applied with different combinations of distillation layers.

B0	B1	B2	H	mAP	NDS
				32.8	45.9
			✓	34.9	46.2
		✓	✓	36.1	47.3
	✓	✓	✓	36.8	48.4
✓	✓	✓	✓	34.4	46.6

(c) Evaluation of the distillation effects at different layers.

Method	FG&BG	Attention	Scaling	FP	mAP	NDS
Baseline					32.8	45.9
Distill					32.9	45.6
	✓				33.4	46.3
	✓	✓			34.6	46.5
	✓	✓		✓	36.8	48.4
	✓	✓	✓	✓	37.8	49.0

(d) Comparison of our approach using different combinations of the proposed region decomposition mask including regions from FG&BG and FP, spatial attention map, as well as adaptive scaling factor.

Layer	Adaptation Module	mAP	NDS
Pre-Head (H)	Linear	37.6	48.8
	2×Block	37.8	49.0
	3×Block	37.0	48.6
Intermediate (B2-B1)	Linear	37.3	48.9
	Upsample-2×Block	37.0	48.8
	Upsample-3×Block	37.8	49.0
	Downsample-3×Block	37.3	48.1

(e) Comparison of different design choices of the adaptation module at different distillation layers. A linear projection is a 1×1 Conv and each block consists of 1×1 Conv-BN-ReLU.

Table 4. A set of ablation study and related design choices on the validation set of nuScenes.

composition for a fair comparison among different levels. For the cross-modal distillation, the teacher and student networks are separately developed with distinct architectures, and the modality gap further complicates the problem. This differs from 2D object detection distillation [13, 44], where the teacher and student models can adopt similar architectures and both take images as input. To alleviate this issue, we align features at similar abstraction level via the adaptation module. We use H to indicate the pre-head layer, and B2-B0 to denote its three preceding intermediate layers of the BEV encoders in both teacher and student. As shown in Table 4c, our multi-scale distillation shows a clear advantage over distilling at the pre-head layer alone (36.8% vs 34.9% mAP, 48.4% vs 46.2% NDS). It is also noteworthy that aligning features from an early layer (B0) is negative, presumably because the representation gap between point clouds and images remains large at the early stage.

False Positive Region Decomposition. As demonstrated in Table 4d, FP regions decomposed from background are beneficial to distillation. Built upon the optimal multi-scale setup in Table 4c, we evaluate different FP region mining principles according to (2), namely the teacher-generated or student-generated FP regions under different thresholds γ . As shown in Table 4a, the teacher-generated FP regions with $\gamma = 0.1$ overall performs the best. This shows that the feature responses even at FP regions produced by teacher is assistive in 3D geometric knowledge transfer. We further evaluate this strategy on multi-scale layers. As compared in Table 4b, FP region decomposition performs better with the pre-head layer only. We hypothesize that such regions can be better represented by the features with high-level semantics that emerge at the last layer.

Adaptation Module. A lightweight adaptation module is employed to map the student feature before aligning to the

teacher feature at corresponding level. This operation is necessary to match the feature size (spatial and channel dimensions) for multi-scale distillation between two discrepant architectures, and to provide extra flexibility to mitigate the difficulty in cross-modal feature imitation learning. We evaluate different design choices of the adaptation module in Table 4e. It is observed that a simple linear projection works comparably well for the pre-head layer, while the intermediate layers require more non-linear mapping. This comparison suggests that the student features from intermediate layers need more adaptation in order to better align with the teacher features due to the relatively large representation gap at the early stage caused by different modalities. As for dimension matching, we can either upsample student features or downsample teacher features. Table 4e shows that upsampling student features is favorable as it preserves more spatial details for feature alignment.

Teacher Model. Finally, we investigate the effect of different teacher models for knowledge distillation in BEV. Here we use BEVDet as the student model, and choose three teacher models including BEVDepth, CenterPoint and MVP, which represent the multi-camera BEV, LiDAR and camera-LiDAR fusion based detectors. Table 6 shows that the distillation across modalities clearly outperforms that within camera modality, validating the advantage of features extracted from point clouds to guide the student detector for more effective representation learning. We can further improve the student model by switching to the sensor-fusion based teacher model MVP, which shares similar network capacity as CenterPoint (both are much more compact than BEVDepth). As shown in Table 5, compared to CenterPoint, MVP boosts the overall performance, in particular for small objects such as pedestrian and motorcycle. This is due to the fact that point clouds become increasingly sparse

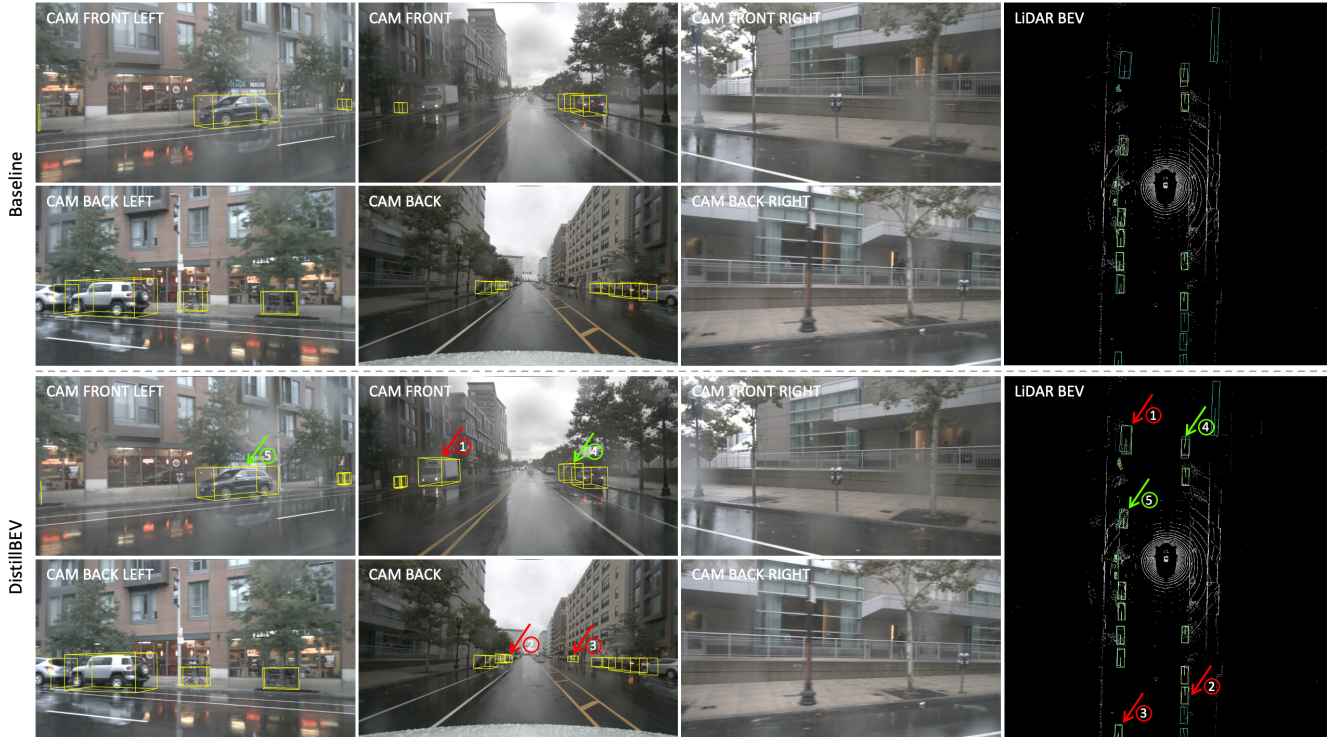


Figure 3. Comparison of the baseline (BEVDepth) and its distilled version by our approach. The cyan (in BEV only for figure clearance) and yellow boxes denote the ground truth and detection results, respectively. We use the red arrows to mark the objects that are missed by baseline but detected by DistillBEV, and the green arrows to indicate the objects that are localized more accurately by DistillBEV.

Teacher	Student	Car	Truck	Bus	Trailer	CV	Ped	Motor	Bicycle	TC	Barrier
BEVDepth	BEVDet	51.4	23.5	34.0	14.9	7.1	33.7	27.8	23.7	49.2	49.5
CenterPoint	BEVDet	54.3	26.8	38.8	19.2	6.6	33.4	27.8	22.0	47.6	50.1
MVP	BEVDet	54.2	28.3	40.4	16.7	7.4	35.2	30.2	25.9	49.7	52.4

Table 5. Comparison of the per-class AP on the validation set of nuScenes based on the three different teacher models to perform knowledge distillation in BEV. Abbreviations are construction vehicle (CV), pedestrian (Ped), motorcycle (Motor), and traffic cone (TC).

Teacher	Mode	Params (M)	mAP	NDS
-	-	-	30.5	37.8
BEVDepth	C	53.2	31.5	38.5
CenterPoint	L	6.0	32.7	40.7
MVP	L&C	6.0	34.0	41.6

Table 6. Comparison of the three different teacher models for knowledge distillation in BEV on the validation set of nuScenes.

for small or distant objects, and augmenting point clouds with dense image information alleviates this issue.

4.4. Qualitative Results

In Figure 3, we visualize the detection results in multi-camera images and BEV (LiDAR is provided for reference). DistillBEV successfully detects three objects that are missed by the baseline, in particular for the two distant vehicles indexed by “2” and “3”. One can also see that our approach localizes objects with more accurate depth as

marked by the green arrows. These examples qualitatively show the efficacy of the proposed approach to enhance the multi-camera BEV based 3D object detectors.

5. Conclusion

We have presented DistillBEV, a simple and effective approach that guides the learning process of a multi-camera BEV based student model through a more powerful LiDAR based teacher model. Extensive experiments on eight teacher-student combinations reveal that our approach consistently renders significant performance gains with no extra computation cost during inference. This demonstrates that our approach is not just geared to a particular network but is flexible to cope with a broad range of models in both CNNs and Transformers. In the future work, we intend to explore and apply the proposed approach to more multi-camera perception tasks in BEV, such as segmentation, tracking and online high-definition map construction.

References

- [1] Tesla AI Day. https://youtu.be/ODSJsviD_SU, 2022.
- [2] Eli Bronstein, Mark Palatucci, Dominik Notz, Brandyn White, Alex Kuefler, Yiren Lu, Supratik Paul, Payam Nikdel, Paul Mouglin, Hongge Chen, Justin Fu, Austin Abrams, Punit Shah, Evan Racah, Benjamin Frenkel, Shimon Whiteson, and Dragomir Anguelov. Hierarchical model-based imitation learning for planning in autonomous driving. In *IROS*, 2022.
- [3] 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 *CVPR*, 2020.
- [4] Yigit Baran Can, Alexander Liniger, Danda Pani Paudel, and Luc Van Gool. Structured bird’s-eye-view traffic scene understanding from onboard images. In *ICCV*, 2021.
- [5] Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qin-hong Jiang, and Feng Zhao. BEVDistill: Cross-modal BEV distillation for multi-view 3D object detection. In *ICLR*, 2023.
- [6] Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, and Wanli Ouyang. MonoDistill: Learning spatial features for monocular 3D object detection. In *ICLR*, 2022.
- [7] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *CVPR*, 2021.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [9] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv:1503.02531*, 2015.
- [10] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [11] Junjie Huang and Guan Huang. BEVDet4D: Exploit temporal cues in multi-camera 3D object detection. *arXiv:2203.17054*, 2022.
- [12] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. BEVDet: High-performance multi-camera 3D object detection in bird-eye-view. *arXiv:2112.11790*, 2021.
- [13] Zijian Kang, Peizhen Zhang, Xiangyu Zhang, Jian Sun, and Nanning Zheng. Instance-conditional knowledge distillation for object detection. In *NeurIPS*, 2021.
- [14] Nikolaus Kriegeskorte. Deep neural networks: A new framework for modelling biological vision and brain information processing. *Annual Review of Vision Science*, 2015.
- [15] Alex Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast encoders for object detection from point clouds. In *CVPR*, 2019.
- [16] Youngwan Lee, Joong-won Hwang, Sangrok Lee, Yuseok Bae, and Jongyoul Park. An energy and GPU-computation efficient backbone network for real-time object detection. In *CVPR Workshop*, 2019.
- [17] Jinyu Li, Chenxu Luo, and Xiaodong Yang. PillarNeXt: Rethinking network designs for 3D object detection in LiDAR point clouds. In *CVPR*, 2023.
- [18] Weixin Li and Xiaodong Yang. Transcendental idealism of planner: Evaluating perception from planning perspective for autonomous driving. In *ICML*, 2023.
- [19] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with Transformer for 3D object detection. *NeurIPS*, 2022.
- [20] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. BEVDepth: Acquisition of reliable depth for multi-view 3D object detection. In *AAAI*, 2023.
- [21] 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. In *ECCV*, 2022.
- [22] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [23] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. PETR: Position embedding transformation for multi-view 3D object detection. In *ECCV*, 2022.
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- [25] Chenxu Luo, Xiaodong Yang, and Alan Yuille. Exploring simple 3D multi-object tracking for autonomous driving. In *ICCV*, 2021.
- [26] Chenxu Luo, Xiaodong Yang, and Alan Yuille. Self-supervised pillar motion learning for autonomous driving. In *CVPR*, 2021.
- [27] Xinzhu Ma, Yinmin Zhang, Dan Xu, Dongzhan Zhou, Shuai Yi, Haojie Li, and Wanli Ouyang. Delving into localization errors for monocular 3D object detection. In *CVPR*, 2021.
- [28] Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled Rezaat, and Benjamin Sapp. Wayformer: Motion forecasting via simple and efficient attention networks. In *ICRA*, 2023.
- [29] Bowen Pan, Jiankai Sun, Ho Yin Tiga Leung, Alex Andonian, and Bolei Zhou. Cross-view semantic segmentation for sensing surroundings. *IEEE Robotics and Automation Letters*, 2020.
- [30] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-LiDAR needed for monocular 3D object detection? In *ICCV*, 2021.
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, and Luca Antiga. PyTorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32, 2019.
- [32] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3D. In *ECCV*, 2020.
- [33] Charles Qi, Hao Su, Kaichun Mo, and Leonidas Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2018.
- [34] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for thin deep nets. *arXiv:1412.6550*, 2014.

- [35] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, 2019.
- [36] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. FCOS3D: Fully convolutional one-stage monocular 3D object detection. In *ICCV*, 2021.
- [37] Xishun Wang, Tong Su, Fang Da, and Xiaodong Yang. ProphNet: Efficient agent-centric motion forecasting with anchor-informed proposals. In *CVPR*, 2023.
- [38] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. DETR3D: 3D object detection from multi-view images via 3D-to-2D queries. In *CoRL*, 2022.
- [39] Yue Wang and Justin Solomon. Object DGCNN: 3D object detection using dynamic graphs. *NeurIPS*, 2021.
- [40] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely embedded convolutional detection. *Sensors*, 2018.
- [41] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. In *CVPR*, 2022.
- [42] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3D object detection and tracking. In *CVPR*, 2021.
- [43] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multi-modal virtual point 3D detection. *NeurIPS*, 2021.
- [44] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *ICLR*, 2021.
- [45] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3D object detection. In *CVPR*, 2018.
- [46] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3D object detection. *arXiv:1908.09492*, 2019.