






LabelDistill: Label-guided Cross-modal Knowledge Distillation for Camera-based 3D Object Detection

Sanmin Kim¹, Youngseok Kim^{2*}, Sihwan Hwang¹, Hyeonjun Jeong¹,
and Dongsuk Kum¹

¹ Korea Advanced Institute of Science and Technology, Daejeon, South Korea
{sanmin.kim, shhwang0129, hyeonjun.jeong, dskum}@kaist.ac.kr

² 42dot Inc., Seoul, South Korea
youngseok.kim@42dot.ai

Abstract. Recent advancements in camera-based 3D object detection have introduced cross-modal knowledge distillation to bridge the performance gap with LiDAR 3D detectors, leveraging the precise geometric information in LiDAR point clouds. However, existing cross-modal knowledge distillation methods tend to overlook the inherent imperfections of LiDAR, such as the ambiguity of measurements on distant or occluded objects, which should not be transferred to the image detector. To mitigate these imperfections in LiDAR teacher, we propose a novel method that leverages aleatoric uncertainty-free features from ground truth labels. In contrast to conventional label guidance approaches, we approximate the inverse function of the teacher’s head to effectively embed label inputs into feature space. This approach provides additional accurate guidance alongside LiDAR teacher, thereby boosting the performance of the image detector. Additionally, we introduce feature partitioning, which effectively transfers knowledge from the teacher modality while preserving the distinctive features of the student, thereby maximizing the potential of both modalities. Experimental results demonstrate that our approach improves mAP and NDS by 5.1 points and 4.9 points compared to the baseline model, proving the effectiveness of our approach. The code is available at <https://github.com/sanmin0312/LabelDistill>

Keywords: Multi-view 3D object detection · Knowledge distillation

1 Introduction

3D object detection is an essential task in various applications, such as autonomous driving and robotics. In recent years, camera-based methods [35, 44, 55, 56] have attracted extensive attention owing to their cost-effectiveness and rich semantic information that images can provide. However, their current performance falls short when compared to LiDAR-based counterparts [26, 57, 63], primarily due to the absence of geometric and spatial information.

*Work done at Korea Advanced Institute of Science and Technology.




Fig. 1: (a) **Conventional cross-modal knowledge distillation** trains an image detector to mimic the features of a well-trained LiDAR detector. It could be suboptimal as it directly transfers LiDAR features with inherent imperfections to the image feature. (b) **LabelDistill** enhances the image detector by incorporating ground truth labels into the feature representation. This approach aims to furnish the image detector with more accurate guidance, alleviating the intrinsic limitations of LiDAR point clouds.

To bridge this performance gap between the camera and LiDAR detectors, knowledge distillation [15] emerges as a promising solution, following the success in various computer vision fields such as image classification [64], object detection [62] and segmentation [34]. Notably, LiDAR-guided cross-modal knowledge distillation methods [6, 8, 16, 19, 22, 25, 27, 59] hold great potential in the camera-based 3D object detection task. These methods transfer learned information from LiDAR detectors to image detectors, leveraging precise spatial features from LiDAR without requiring LiDAR sensors during inference.

Despite the improvements observed in current LiDAR-guided cross-modal knowledge distillation methods, they are not without their limitations. First, they tend to overlook the inherent imperfections of LiDAR point clouds, including aleatoric uncertainties in distant and occluded objects. Such shortcomings make features from LiDAR detector imperfect for distillation. Second, existing methods insufficiently handle complementary characteristics of LiDAR and camera. While LiDAR provides precise spatial information, the camera offers abundant semantic information. Therefore, indiscriminate distillation aiming to align all image features with LiDAR features may hinder the extraction of the full potential of image features.

To address these limitations, we present a novel cross-modal knowledge distillation approach tailored for camera-based 3D object detection. Our approach introduces a label distillation strategy that capitalizes on aleatoric uncertainty-free features derived from ground truth labels within the distillation process. Unlike conventional label guidance approaches [14, 68], which extract label features supervised by student features, our label distillation method focuses on extracting label features that can complement the limitations of LiDAR point clouds. This is achieved by leveraging the inverse function of a well-trained teacher’s head, which can effectively map 3D bounding boxes into a teacher’s feature space. When combined with LiDAR distillation, our label distillation approach provides accurate and robust guidance to the image detector, enhancing its overall performance.

Furthermore, we introduce a feature partitioning strategy in the distillation process to effectively transfer knowledge from the teacher modality while preserving the complementary features of the student modality, such as semantic information. We separate student’s features into several groups in the channel dimension, allocating some to the teacher while keeping others unaffected by the teacher. This approach ensures that the student can learn informative features from the teacher without compromising its own unique characteristics. In summary, the contributions of this paper are:

- We propose a novel label-guided cross-modal knowledge distillation, which effectively complements the imperfections of the LiDAR-based teacher model, leveraging the aleatoric uncertainty-free features.
- We introduce a feature partitioning to effectively transfer knowledge from the teacher modality while preserving the distinctive information of the student modality.
- Our approach achieves improved performance compared to prior state-of-the-art methods without incurring additional costs in the inference stage. Extensive experimentation confirms the effectiveness of our approach.

2 Related Work

Camera-based 3D Object Detection. Early approaches in camera-based 3D object detection [1, 37, 42, 44, 55] built upon the success of 2D detection methods [51, 72]. These methods utilized perspective view features to directly estimate 3D information from 2D image inputs. However, they faced the challenge of ill-posed depth estimation, stemming from information loss during the projection from 3D to 2D. To mitigate such inaccurate depth estimation, several methods [32, 40, 48, 54] have explored geometric information, while DD3D [44] have incorporated depth pre-training using additional datasets [11].

Recent progress in the field have involved the adoption of of Bird’s-Eye-View (BEV) feature representation through view transformation. A line of works [18, 29, 47, 49] has adopted forward view transformations by projecting perspective view features into BEV space using estimated depth distribution. On the other hand, other works [5, 23, 30, 50, 56, 61] have employed backward view transformation by incorporating attention mechanism [53] for correspondences between 3D and 2D space. Despite these advancements in camera-based 3D object detection showing promising performance, challenges persist in achieving accurate localization due to the inherent limitations of depth information.

Knowledge Distillation for 3D Object Detection. Knowledge distillation is initially proposed for the model compression [15] by transferring the information from a large and cumbersome teacher model to a light and compact student model. It has proven effective in various computer vision domains, such as classification [46, 60, 64], object detection [3, 9, 62], and semantic segmentation [34, 52, 58]. Recently, this strategy has been applied to the 3D object detection task [7, 65, 67, 69].

In autonomous driving applications, LiDAR-guided cross-modal knowledge distillation methods [6, 8, 16, 19, 22, 25, 27, 59, 71] are gaining attention, which introduce a LiDAR detector as the teacher model to provide accurate and rich spatial information obtained from LiDAR point clouds to an image detector. MonoDistill [8] projects LiDAR points into the image plane to unify the representations, and BEVDistill [6] introduces a sparse instance-wise distillation in addition to dense feature imitation. On the other hand, X³KD [25] proposes cross-task knowledge distillation that transfers information from instance segmentation tasks. Despite their promising results, these methods often overlook the imperfections in LiDAR data, leading to suboptimal distillation. Additionally, domain discrepancies between LiDAR and camera modalities are insufficiently addressed.

Label Guidance. Several works across various tasks have integrated label guidance into their training schemes. One line of work [41, 43] employs labels for intermediate supervision, offering auxiliary guidance for regularization. Another line of work [14, 21, 38, 68] utilizes label input to enhance student features within a teacher-free distillation framework. However, these methods struggle to effectively extract useful features from labels as they typically employ simplistic autoencoders or rely on student features to train the label encoder, resulting in suboptimal label features. In contrast, our approach involves embedding labels into the feature space of a LiDAR teacher model, thereby providing valuable label features that can complement teacher features.

3 Method

As illustrated in Fig. 2, our proposed method consists of three pipelines: LiDAR, ground truth labels, and image. The primary goal is to guide the image detector in learning accurate spatial information by employing label distillation in addition to LiDAR distillation, all while preserving its distinctive features.

3.1 LiDAR Distillation

The LiDAR distillation process follows the conventional knowledge distillation paradigm, utilizing a LiDAR detector as the teacher model. Our approach begins by extracting Bird’s-Eye-View (BEV) features from both LiDAR point clouds and multi-view images, employing independent backbones for each modality. We utilize two LiDAR distillation strategies: feature-level and response-level distillation.

Feature-level Distillation. Feature-level distillation aims to transfer rich spatial and geometric information from LiDAR BEV features to the corresponding image BEV features. These image BEV features are transformed from the perspective view using view transformation techniques [30, 47]. This distillation is facilitated through a loss function as follows:

$$\mathcal{L}_{lidar}^{feat} = \frac{1}{N_p} \sum_i^H \sum_j^W \mathcal{M}_{ij} \{F_{ij}^{lidar} - \alpha(F_{ij}^{image})\}^2, \quad (1)$$




Fig. 2: Overall architecture of the proposed method. Our model is trained with two distillation strategies: LiDAR distillation and label distillation. **LiDAR Distillation** transfers abundant spatial information to the image detector using feature-level and response-level distillation. **Label Distillation** provides accurate and aleatoric uncertainty-free information based on the ground truth label to compensate the limitations of LiDAR point clouds. In addition, **Feature Partitioning** separates the image features into three groups to preserve distinctive image features while learning from LiDAR and label features.

where H and W represent the height and width of the BEV feature map. F_{ij}^{lidar} and F_{ij}^{image} are the BEV features at location (i, j) from the LiDAR and image, respectively. The mask \mathcal{M} isolates the distillation process to object-specific regions, employing a foreground mask derived from the ground truth heatmap within the BEV space. N_p is the number of non-zero pixels in \mathcal{M} . The adaptation module α , consisting of convolutional layers, aligns the dimensionality of the image features with the teacher model’s output.

Response-level Distillation. In response-level distillation, the predictions from the LiDAR detector are used as an additional soft label, following [15]:

$$\mathcal{L}_{lidar}^{resp} = \mathcal{L}_{cls}(c_{lidar}, c_{image}) + \mathcal{L}_{bbox}(b_{lidar}, b_{image}), \quad (2)$$

where c and b denote the class heatmap and bounding box predictions from LiDAR and image detector, respectively. We employ focal loss for the classification loss \mathcal{L}_{cls} and L1 loss for the regression loss \mathcal{L}_{bbox} . In this process, we utilize foreground masking based on ground truth heatmaps to prevent negative impacts from false positives.

3.2 Label Distillation

While LiDAR distillation provides essential spatial information to guide the image detector, the inherent limitations of LiDAR point clouds, such as ambiguity in distant or occluded objects due to sparsity [66] and susceptibility to adverse weather [12, 13], can potentially impact the quality of features used in the distillation process. These imperfections tend to be neglected in existing

studies since they were overshadowed by the superior detection performance of the LiDAR object detectors over camera detectors, thereby limiting the full potential of LiDAR-guided cross-modal knowledge distillation. To overcome these limitations, we introduce label distillation as a complementary strategy alongside LiDAR distillation. The label distillation leverages the ground truth labels. The ground truth labels are generated by human annotators using multiple sensors with long sequential frames (*e.g.*, the nuScenes dataset [2] leverages LiDAR, radar, and camera with 20 seconds of frames, including past and future timesteps). As a result, these ground truth labels are ready to offer precise 3D object bounding boxes that are free from aleatoric uncertainty, providing the image detector with reliable guidance.

Approximating the Inverse Function of the Teacher’s Head. A crucial step in mitigating the limitations of the teacher model is to adequately encode ground truth labels into the feature space. Previous efforts to utilize labels for guiding the training process have been explored in several works [14, 21, 68]. However, these methods have often fallen short in extracting optimal features from label inputs, primarily due to a training process that forces label features to be similar to student features. To address this challenge, we leverage the LiDAR detection head’s capability that decode LiDAR features into 3D bounding box predictions:

$$\hat{y} = h(F_{lidar}; \theta_h), \quad (3)$$

where F_{lidar} and \hat{y} denote LiDAR features and the bounding box predictions, respectively. $h(\cdot; \theta_h)$ represents the LiDAR detection head.

This process implies that the inverse function of the LiDAR detection head can map bounding box representations back into feature space. Accordingly, we aim to embed labels, which are 3D bounding boxes, into the feature space of the teacher model using this inverse function of the LiDAR detection head, as formalized in the following equation:

$$F_{label} = h^{-1}(y; \theta_{h^{-1}}), \quad (4)$$

where $h^{-1}(\cdot; \theta_{h^{-1}})$ represents the inverse function of the LiDAR detection head, acting as the label encoder. In other words, h^{-1} can output optimal label features given ground truth 3D bounding box inputs.

However, calculating this inverse function is impractical due to the high non-linearity of neural networks. Inspired by [14] and [41], we utilize an autoencoder framework to approximate the inverse function of the LiDAR detection head. Within this framework, the label encoder assumes the role of the encoder, and the pre-trained LiDAR detection head functions as the decoder, as described in Fig. 3. The training objective for the label encoder is formulated as:

$$\theta_g^* = \arg \min_{\theta_g} \mathbb{E}_{(I, y) \sim \mathcal{D}} \mathcal{L}_{det} \left(h(g(y; \theta_g); \theta_h^*), y \right), \quad (5)$$

where $h(\cdot; \theta_h^*)$ represents the pretrained LiDAR detection head, and $g(\cdot; \theta_g)$ represents the label encoder designed to approximate $h^{-1}(\cdot; \theta_h^*)$. (I, y) denotes a pair of LiDAR point cloud and ground truth label, \mathcal{D} is the distribution of the dataset,




Fig. 3: Architecture of the label encoder. The label encoder is designed to approximate the inverse function of the pretrained lidar detection head. The label encoder first encodes class and bounding box information and then, the mapping function transforms encoded label features into BEV space by filling the object’s bounding box area with label features. Finally, the convolutional block encodes BEV label features.

and $\mathcal{L}_{det}(\cdot, \cdot)$ represents the detection loss function for the classification and bounding box regression. In this manner, label distillation effectively mitigates imperfections of LiDAR point clouds by aligning the aleatoric uncertainty-free label features to the teacher’s feature space. Our approach differs from conventional autoencoders by setting the decoder as the pretrained LiDAR detection head and focusing the training on the encoder (the label encoder), whereas a standard autoencoder would train both components from scratch.

Label Encoder. As depicted in Fig. 3, we have adopted a simple design for the label encoder due to the compact and noise-free nature of ground truth labels. The label encoder handles both class and bounding box information, employing a straightforward yet efficient structure for label encoding. The label encoder is defined as follows:

$$g(y; \theta_g) = f\left(q\left(\Phi_{cls}(c_{gt}) + \Phi_{box}(b_{gt})\right)\right), \quad (6)$$

where $c_{gt} \in \mathbb{R}^{n \times m}$ represents the ground truth class information of m classes for n objects in the scene, while $b_{gt} \in \mathbb{R}^{n \times z}$ is the ground truth bounding box information of z attributes such as 3D location, size, orientation and velocity. Φ_{cls} and Φ_{box} are MLP layers to embed class and bounding box information. The embedded class and bounding box vectors are placed in foreground on the BEV space using the mapping function $q(\cdot)$ after summation. We fill each BEV grid occupied by the bounding box of an object, with duplicated label feature vectors, to generate the label BEV feature. Subsequently, the function f , which is the convolutional block including convolutional layers, normalization, and activations, is employed to refine the feature maps into the final label feature F_{label} . Note that the label encoder is pretrained before the distillation process.

The implementation of this label encoder has proven to be highly effective in approximating the inverse function of the LiDAR detection head, achieving a

Table 1: Evaluation of the autoencoder consists of the label encoder and LiDAR detection head on nuScenes validation set.

	mAP \uparrow	NDS \uparrow	mATE \downarrow	mAOE \downarrow	mAVE \downarrow
Label Encoder + LiDAR Head	94.14	90.25	0.192	0.048	0.128

94% mean Average Precision (mAP) when combined with the LiDAR detection head, as illustrated in Tab. 1. This result also demonstrates that the encoded label features preserve useful information to reconstruct 3D bounding boxes while mapping to the teacher’s feature space.

3.3 Feature Partitioning

LiDAR point clouds are a rich source of precise spatial and geometric data, while images offer dense semantic details. These modalities are inherently complementary. However, conventional cross-modal distillation approaches that attempt to train all image feature channels to mimic LiDAR features may not fully harness the potential of images. Additionally, our approach utilizes multiple teachers, including LiDAR and label, which can potentially result in contrasting supervision due to the disparate nature of features.

To address these challenges, we introduce a straightforward yet effective strategy: feature partitioning. This strategy aims to preserve distinctive image features while simultaneously learning from the LiDAR and label features. We partition the image feature $F_{\text{image}} \in \mathbb{R}^{H \times W \times C}$ into three distinct groups along the channel dimension: $F_{\text{image}}^{\text{image}}, F_{\text{image}}^{\text{lidar}}, F_{\text{image}}^{\text{label}}$. Each feature group consists of a subset of the image features with a combined total of C channels. The group $F_{\text{image}}^{\text{lidar}}$ is designed to focus on learning essential LiDAR features, leveraging the spatial and geometric details provided by the LiDAR data. Meanwhile, the group $F_{\text{image}}^{\text{label}}$ is dedicated to learning label-related features. In contrast, the group $F_{\text{image}}^{\text{image}}$ remains unaffected by the influence of the teacher models. This group is exclusively trained using the detection loss function. By remaining uninfluenced by the teacher models, this group retains the inherent semantic features found in the image data, ensuring that the richness and depth of the semantic information remain intact throughout the training process.

3.4 Training

Our model undergoes a two-step training process. In the first step, we train the label encoder to approximate the inverse function of the pretrained LiDAR detection head. During this step, the label encoder is trained using a conventional detection loss, including classification and bounding box regression losses. In the second step, we train the image detector with the pretrained label encoder and LiDAR detector. This step involves training our model with a loss function

comprising four terms: LiDAR feature loss, LiDAR response loss, label feature loss, and detection loss, which is formulated as:

$$\mathcal{L} = \mathcal{L}_{det} + \lambda_1 \mathcal{L}_{lidar}^{feat} + \lambda_2 \mathcal{L}_{label}^{feat} + \lambda_3 \mathcal{L}_{lidar}^{resp}, \quad (7)$$

where $\lambda_{1,2,3}$ are balancing weight term. We adopt the same loss function as presented in [29] for the detection loss. It consists of classification loss, bounding box regression loss, and depth loss. Meanwhile, the label feature loss employs the Mean Squared Error (MSE) loss with foreground masking, similar to the LiDAR feature loss in Eq. (1). It is worth to note that our distillation strategies do not introduce any additional computational burden during the inference stage.

4 Experiments

4.1 Experimental Setup

Dataset and Metrics. We train and evaluate our approach on the nuScenes dataset [2], which is the large-scale autonomous driving benchmark. It consists of 1000 videos of around 20 seconds with annotations at 2Hz, including 3D bounding boxes of 10 classes. We follow the official evaluation metrics to evaluate 3D object detection performance, including mean Average Precision (mAP) and nuScenes Detection Score (NDS). We also report other metrics such as mean Average Translation Error (mATE), mean Average Scale Error (mASE), mean Average Orientation Error (mAOE), mean Average Velocity Error (mAVE), and mean Average Attribute Error (mAAE).

Teacher and Student Model. For teacher model, we adopt pretrained CenterPoint [63] with a voxel size of (0.1m, 0.1m, 0.2m). For student model, we employ BEVDepth [29]. Unless otherwise specified, ResNet50 pretrained with ImageNet is adopted as image backbone, and the input image is resized to 256×704 . We follow the image and BEV data augmentation strategies in [29]. We use four previous frames for the experiments of Tab. 2 and Tab. 3 while one previous frame is adopted for ablation studies.

Implementation Details. The label encoder is trained for 12 epochs with the learning rate of 1e-3 while 24 epochs and learning rate of 4e-4 is employed for training the image detector. We adopt AdamW optimizer [39] without CBGS [73]. A batch size of 16 on 4 NVIDIA 3090Ti GPUs is used for both training of label encoder and distillation of student model.

4.2 Main Results

We start our analysis by comparing our model with existing camera-based 3D object detection models on the nuScenes validation set. As reported in Tab. 2, our model achieves a significant improvement of 8.6%p in mAP and 8.7%p in NDS compared to the baseline model, BEVDepth, in the ResNet50 settings. Notably, these improvements remain consistent with a 4.5%p and 6.3%p boost in

Table 2: Comparison on the nuScenes dataset. †: methods with CBGS. *: we reproduced BEVDepth with the same setting as our model for a fair comparison.

Set	Method	Backbone	Size	mAP	NDS	mATE	mASE	mAOE	mAVE	mAAE
Validation	BEVDet4D [17]	ResNet50	256×704	32.3	45.3	0.674	0.272	0.503	0.429	0.208
	BEVDepth [29]	ResNet50	256×704	33.3	44.1	0.683	0.276	0.545	0.526	0.226
	BEVStereo [28]	ResNet50	256×704	34.4	44.9	0.659	0.276	0.579	0.503	0.216
	VEDet† [4]	ResNet50	384×1056	34.7	44.3	0.726	0.282	0.542	0.555	0.198
	PETR v2 [36]	ResNet50	256×704	34.9	45.6	0.700	0.275	0.580	0.437	0.187
	FB-BEV† [31]	ResNet50	256×704	35.0	47.9	0.642	0.275	0.459	0.391	0.193
	AeDet† [10]	ResNet50	256×704	35.8	47.3	0.655	0.273	0.493	0.427	0.216
	P2D [24]	ResNet50	256×704	37.4	48.6	0.631	0.272	0.508	0.384	0.212
	BEVFormer v2† [61]	ResNet50	640×1600	38.8	49.8	0.679	0.276	0.417	0.403	0.189
	SOLOFusion [45]	ResNet50	256×704	40.6	49.7	0.609	0.284	0.650	0.315	0.204
	LabelDistill	ResNet50	256×704	41.9	52.8	0.582	0.258	0.413	0.346	0.220
Validation	DETR3D† [56]	ResNet101	900×1600	34.9	43.4	0.716	0.268	0.379	0.842	0.200
	BEVDepth [29]	ResNet101	512×1408	40.6	49.0	0.626	0.278	0.513	0.489	0.226
	BEVFormer [30]	ResNet101	900×1600	41.6	51.7	0.673	0.274	0.372	0.394	0.198
	VEDet† [4]	ResNet101	512×1408	43.2	52.0	0.638	0.275	0.362	0.498	0.191
	PolarFormer [23]	ResNet101	900×1600	43.2	52.8	0.648	0.270	0.348	0.409	0.201
	P2D [24]	ResNet101	512×1408	43.3	52.8	0.619	0.265	0.432	0.364	0.211
	Sparse4D [33]	ResNet101	900×1600	43.6	54.1	0.633	0.279	0.363	0.317	0.177
	LabelDistill	ResNet101	512×1408	45.1	55.3	0.579	0.252	0.331	0.357	0.207
Test	BEVDepth* [29]	ConvNeXt-B	900×1600	47.5	56.1	0.474	0.259	0.463	0.432	0.134
	LabelDistill	ConvNeXt-B	900×1600	52.6	61.0	0.443	0.241	0.339	0.370	0.136

mAP and NDS, respectively, even in the ResNet101. Furthermore, it also demonstrates superior performance compared to other state-of-the-art approaches. It is noteworthy that our model attains these results without resorting to CBGS [73], a data augmentation strategy that effectively extends a single epoch into 4.5 epochs.

In addition, we perform a comparative analysis of our model with other LiDAR-guided cross-modal knowledge distillation methods, as shown in Table Tab. 3. For a fair comparison, we present performance gain from baselines (Δ) for both mAP and NDS. This metric allows for a simple and equitable comparison as each model shares the same experimental settings with its baseline. As shown in Tab. 3, our approach achieves superior performance compared to these models.

In the case of the test set, we trained BEVDepth [29] with the same settings as our model to ensure fair comparison. As a results, our LabelDistill achieves improvement of 5.1%p and 4.9%p for mAP and NDS, respectively.

4.3 Ablation Study

We performed a series of comprehensive ablation studies to evaluate the contribution of individual components and the impact of different hyperparameters within our model. These studies were conducted on the nuScenes validation set, with results detailed in Tab. 4 through Tab. 8.

Table 3: Comparison to other LiDAR-guided cross-modal knowledge distillation strategies. †: methods with CBGS.

Model	Baseline	Image Size	Backbone	mAP (Δ)	NDS (Δ)
UniDistill [71]	BEVDet	704×256	ResNet50	29.6 (3.2)	39.3 (3.2)
BEVDistill [6]	BEVDepth	704×256	ResNet50	33.0 (1.3)	45.2 (1.2)
TiG-BEV [20]	BEVDepth	704×256	ResNet50	36.6 (3.7)	46.1 (3.0)
BEVSimDet [70]	BEVFusion-C	704×256	ResNet50	37.3 (1.7)	43.8 (2.6)
X ³ KD [†] [25]	BEVDepth	704×256	ResNet50	39.0 (3.1)	50.5 (3.3)
DistillBEV [†] [59]	BEVDepth	704×256	ResNet50	40.3 (3.9)	51.0 (2.6)
LabelDistill	BEVDepth	704×256	ResNet50	41.9 (5.1)	52.8 (4.5)
UVTR [27]	-	1600×900	ResNet101	39.2 (1.3)	48.8 (0.5)
BEVDistill [†] [6]	BEVFormer	1600×900	ResNet101	41.7 (1.2)	52.4 (1.8)
TiG-BEV [20]	BEVDepth	1408×512	ResNet101	43.0 (2.4)	51.4 (2.3)
DistillBEV [†] [59]	BEVDepth	1408×512	ResNet101	45.0 (2.3)	54.7 (3.1)
LabelDistill	BEVDepth	1408×512	ResNet101	45.1 (2.4)	55.3 (3.7)

Table 4: Ablation study on the proposed method. LiDAR, Label, and Partition represent LiDAR distillation, label distillation, and feature partitioning, respectively.

	LiDAR	Label	Partition	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow
(a)				33.6	44.8	0.694	0.273
(b)	✓			35.4	48.6	0.648	0.262
(c)	✓	✓		37.0	49.5	0.663	0.258
(d)	✓	✓	✓	37.9	50.1	0.641	0.256

Label Distillation. The ablation comparison presented in Tab. 4 provides a analysis of the effectiveness of each strategy within our proposed model. LiDAR distillation (b) demonstrates improvement in performance compared to the baseline model. However, the integration of label distillation alongside LiDAR distillation (c) yields further enhancement, highlighting the capacity of label distillation to address the limitations of the LiDAR teacher model. Moreover, we offer visual insights into the effectiveness of label distillation through the visualization of Bird’s Eye View (BEV) features, as illustrated in Fig. 4. As depicted in Fig. 4, the label-distilled student feature (F_{image}^{label}) exhibits clear activation, whereas the lidar-distilled student feature (F_{image}^{lidar}) displays either blurry or negligible activation for occluded or distant objects. This observation underscores the superior capability of label distillation in capturing crucial information for challenging scenarios where LiDAR-based features may fall short.

Feature Partitioning. The significance of feature partitioning is underscored by the comparison between (c) and (d) as depicted in Tab. 4. This comparison highlights the advantages conferred by feature partitioning within the distillation process, reaffirming its role in preserving the distinctive image features.

Channel Ratio. We explore the impact of channel ratios on the feature partitioning strategy, maintaining a constant channel ratio for image features while




Fig. 4: Illustration of BEV feature maps in the inference stage. F_{image}^{image} is undistilled image feature, F_{image}^{lidar} is lidar-distilled image feature, and F_{image}^{label} , label-distilled image feature, and F_{label} denotes label feature from the label encoder.

Table 5: Experiments on different channel ratio for the feature partitioning.

Channel Ratio			mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow
F_{lidar}^{image}	F_{label}^{image}	F_{image}^{image}				
1	3	2	36.6	48.8	0.655	0.260
3	1	2	37.1	49.4	0.646	0.258
2	2	2	37.6	49.6	0.643	0.256

varying the ratios for LiDAR and label features. We adopt 300 as the total channels, and as indicated in Tab. 5, the most balanced performance is achieved when the channel ratios for all three features are identical.

Inverse Function Approximation. To evaluate the effectiveness of training the label encoder by approximating the inverse function of the LiDAR detection head, we compared it with other label guidance methods, as shown in Tab. 6. AutoEncoder represents a simplistic approach where both an encoder and decoder are trained from scratch using labels as both inputs and targets. Similarly, LabelEnc [14] employs an AutoEncoder but integrates an additional encoding strategy that relies on the student feature during the label feature training process. In contrast, our method leverages the inverse function of the teacher’s head to effectively embed label features into the feature space. As demonstrated in Tab. 6, our approach outperforms other label guidance methods. These results underscore the effectiveness of employing the inverse function approximation of the teacher head, which ensures accurate and noise-free features are provided to the student model during the distillation process.

Impact of Label Encoder Performance. We examined the influence of the label encoder’s performance on the distillation process. By deliberately reducing the label encoder’s detection capabilities during the label encoder training, we observed a positive correlation between the label encoder’s performance and that of the distilled student model, as shown in Tab. 7. This observation underscores the significance of an accurate inverse function approximation of the teacher detection head in providing precise label features in our distillation strategy.

Table 6: Evaluation on the effectiveness of the inverse function approximation. AutoEncoder trains the label encoder and the detection head from the scratch.

Label Encoder Training	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow
AutoEncoder	34.9	46.7	0.656	0.270	0.476
LabelEnc [14]	34.8	46.8	0.658	0.267	0.479
Inverse Function Approximation	36.8	48.1	0.646	0.263	0.474

Table 7: Experiments of the label encoder’s impact on the student model. Performance of the label encoder denotes AutoEncoder’s performance, which consists of the label encoder and the LiDAR detection head.

Label Encoder + LiDAR Head		Student Model				
mAP \uparrow	NDS \uparrow	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow
50.2	42.9	34.0	45.3	0.678	0.273	0.587
71.9	54.7	34.6	45.6	0.673	0.274	0.583
94.1	90.3	36.8	48.1	0.660	0.264	0.470

Table 8: Performance along the object distance.

Distance	LiDAR	Label	mATE \downarrow	mASE \downarrow	mAOE \downarrow
$\leq 30\text{m}$	\checkmark		0.592	0.261	0.397
	\checkmark	\checkmark	0.582	0.253	0.380
$30\text{m} \leq$	\checkmark		1.043	0.342	0.534
	\checkmark	\checkmark	1.012	0.270	0.531

Distant Objects. An evaluation based on object distance was performed to further explore label distillation’s impact, with findings shown in Tab. 8. The results confirm an overall performance enhancement in models using label distillation. Notably, the size estimation accuracy (mASE) for distant objects (over 30m) is substantially improved when employing label distillation as opposed to solely LiDAR distillation. This improvement can be attributed to the mitigating effect of label distillation on LiDAR sparsity. The sparsity inherent in LiDAR often results in limited points being reflected from distant objects, making size estimation challenging. However, the label distillation resolves this challenge by providing accurate and reliable information, thereby alleviating the impact of sparsity, particularly for distant objects.

4.4 Qualitative Results

In Fig. 5, we visualize a sample case to compare our approach to the baseline model. As indicated with blue circles, LabelDistill demonstrates several key advantages: 1) It achieves higher recall by successfully detecting objects that the baseline model often misses. 2) The accuracy of object localization is notably enhanced. LabelDistill accurately detects the location of objects, whereas the




Fig. 5: Comparison of the baseline (BEVDepth) and our approach. The blue circles in the BEV view highlight cases that demonstrate the advantages of our approach, including: 1) higher recall, 2) more accurate localization, and 3) fewer false positives.

baseline model tends to yield imprecise results. 3) It effectively reduces false positives. LabelDistill reduces unnecessary and redundant bounding boxes while the baseline generates multiple redundant bounding boxes along the depth direction due to its inaccurate depth estimation ability. These advantages make LabelDistill a promising solution for enhancing camera-based 3D object detection in real-world applications.

5 Conclusion

In this paper, we have presented a novel approach for cross-modal knowledge distillation aimed at effectively transferring knowledge from a LiDAR detector to an image detector. Our method, LabelDistill, addresses the inherent imperfections of LiDAR detectors by leveraging precise ground truth labels to provide accurate and aleatoric uncertainty-free features. Additionally, we have introduced a feature partitioning strategy designed to preserve distinctive image features while simultaneously facilitating the learning of accurate spatial information from the teacher model. Our extensive experiments demonstrate the effectiveness of the proposed methods. However, the performance of the proposed method still lags behind compared to those of LiDAR detector.

However, it is important to note that the performance of the proposed method still lags behind compared to those of LiDAR detectors. Furthermore, the effectiveness of our method is dependent on the quality of the ground truth labels. If the ground truth labels in the dataset exhibit low reliability, the performance of the proposed method may be degraded.

Acknowledgements. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) and the National Research Foundation of Korea(NRF) funded by the Korea government(MSIT) under Grants 2021-0-01176 and 2022R1A2C200494413.

References

1. Brazil, G., Liu, X.: M3d-rpn: Monocular 3d region proposal network for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9287–9296 (2019) [3](#)
2. 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) [6](#), [9](#)
3. Cao, W., Zhang, Y., Gao, J., Cheng, A., Cheng, K., Cheng, J.: Pkd: General distillation framework for object detectors via pearson correlation coefficient. *Advances in Neural Information Processing Systems* **35**, 15394–15406 (2022) [3](#)
4. Chen, D., Li, J., Guizilini, V., Ambrus, R.A., Gaidon, A.: Viewpoint equivariance for multi-view 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9213–9222 (2023) [10](#)
5. Chen, S., Wang, X., Cheng, T., Zhang, Q., Huang, C., Liu, W.: Polar parametrization for vision-based surround-view 3d detection. *arXiv preprint arXiv:2206.10965* (2022) [3](#)
6. Chen, Z., Li, Z., Zhang, S., Fang, L., Jiang, Q., Zhao, F.: Bevdistill: Cross-modal bev distillation for multi-view 3d object detection. In: International Conference on Learning Representations (2023) [2](#), [4](#), [11](#)
7. Cho, H., Choi, J., Baek, G., Hwang, W.: itkd: Interchange transfer-based knowledge distillation for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13540–13549 (2023) [3](#)
8. Chong, Z., Ma, X., Zhang, H., Yue, Y., Li, H., Wang, Z., Ouyang, W.: Monodistill: Learning spatial features for monocular 3d object detection. In: International Conference on Learning Representations (2022) [2](#), [4](#)
9. 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) [3](#)
10. Feng, C., Jie, Z., Zhong, Y., Chu, X., Ma, L.: Aedet: Azimuth-invariant multi-view 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21580–21588 (2023) [10](#)
11. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3d packing for self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2485–2494 (2020) [3](#)
12. Hahner, M., Sakaridis, C., Bijelic, M., Heide, F., Yu, F., Dai, D., Van Gool, L.: Lidar snowfall simulation for robust 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16364–16374 (2022) [5](#)
13. Hahner, M., Sakaridis, C., Dai, D., Van Gool, L.: Fog simulation on real lidar point clouds for 3d object detection in adverse weather. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15283–15292 (2021) [5](#)

14. Hao, M., Liu, Y., Zhang, X., Sun, J.: Labelenc: A new intermediate supervision method for object detection. In: European Conference on Computer Vision. pp. 529–545. Springer (2020) [2](#), [4](#), [6](#), [12](#), [13](#)
15. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) [2](#), [3](#), [5](#)
16. Hong, Y., Dai, H., Ding, Y.: Cross-modality knowledge distillation network for monocular 3d object detection. In: European Conference on Computer Vision. pp. 87–104 (2022) [2](#), [4](#)
17. Huang, J., Huang, G.: Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. arXiv preprint arXiv:2203.17054 (2022) [10](#)
18. Huang, J., Huang, G., Zhu, Z., Du, D.: Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. arXiv preprint arXiv:2112.11790 (2021) [3](#)
19. Huang, L., Li, Z., Sima, C., Wang, W., Wang, J., Qiao, Y., Li, H.: Leveraging vision-centric multi-modal expertise for 3d object detection. In: Thirty-seventh Conference on Neural Information Processing Systems (2023) [2](#), [4](#)
20. Huang, P., Liu, L., Zhang, R., Zhang, S., Xu, X., Wang, B., Liu, G.: Tig-bev: Multi-view bev 3d object detection via target inner-geometry learning. arXiv preprint arXiv:2212.13979 (2022) [11](#)
21. Huang, Y., Liu, X., Zhu, Y., Xu, Z., Shen, C., Che, Z., Zhang, G., Peng, Y., Feng, F., Tang, J.: Label-guided auxiliary training improves 3d object detector. In: European Conference on Computer Vision. pp. 684–700 (2022) [4](#), [6](#)
22. Jang, S., Jo, D.U., Hwang, S.J., Lee, D., Ji, D.: Stxd: Structural and temporal cross-modal distillation for multi-view 3d object detection. In: Thirty-seventh Conference on Neural Information Processing Systems (2023) [2](#), [4](#)
23. Jiang, Y., Zhang, L., Miao, Z., Zhu, X., Gao, J., Hu, W., Jiang, Y.G.: Polarformer: Multi-camera 3d object detection with polar transformer. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 1042–1050 (2023) [3](#), [10](#)
24. Kim, S., Kim, Y., Lee, I.J., Kum, D.: Predict to detect: Prediction-guided 3d object detection using sequential images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18057–18066 (2023) [10](#)
25. Klingner, M., Borse, S., Kumar, V.R., Rezaei, B., Narayanan, V., Yogamani, S., Porikli, F.: X3kd: Knowledge distillation across modalities, tasks and stages for multi-camera 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13343–13353 (2023) [2](#), [4](#), [11](#)
26. Koh, J., Lee, J., Lee, Y., Kim, J., Choi, J.W.: Mgtanet: Encoding sequential lidar points using long short-term motion-guided temporal attention for 3d object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 1179–1187 (2023) [1](#)
27. 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 pp. 18442–18455 (2022) [2](#), [4](#), [11](#)
28. 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. pp. 1486–1494 (2023) [10](#)
29. 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. pp. 1477–1485 (2023) [3](#), [9](#), [10](#)
30. 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 (2022) [3](#), [4](#), [10](#)

31. 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) [10](#)
32. Li, Z., Qu, Z., Zhou, Y., Liu, J., Wang, H., Jiang, L.: Diversity matters: Fully exploiting depth clues for reliable monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2791–2800 (2022) [3](#)
33. 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) [10](#)
34. Liu, Y., Chen, K., Liu, C., Qin, Z., Luo, Z., Wang, J.: Structured knowledge distillation for semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2604–2613 (2019) [2](#), [3](#)
35. 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 (2022) [1](#)
36. 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) [10](#)
37. Liu, Z., Wu, Z., Tóth, R.: Smoke: Single-stage monocular 3d object detection via keypoint estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 996–997 (2020) [3](#)
38. Liu, Z., Zhu, L.: Label-guided attention distillation for lane segmentation. *Neurocomputing* **438**, 312–322 (2021) [4](#)
39. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *International Conference on Learning Representations* (2019) [9](#)
40. Lu, Y., Ma, X., Yang, L., Zhang, T., Liu, Y., Chu, Q., Yan, J., Ouyang, W.: Geometry uncertainty projection network for monocular 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3111–3121 (2021) [3](#)
41. Mostajabi, M., Maire, M., Shakhnarovich, G.: Regularizing deep networks by modeling and predicting label structure. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5629–5638 (2018) [4](#), [6](#)
42. Mousavian, A., Anguelov, D., Flynn, J., Kosecka, J.: 3d bounding box estimation using deep learning and geometry. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 7074–7082 (2017) [3](#)
43. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European Conference on Computer Vision. pp. 483–499 (2016) [4](#)
44. 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) [1](#), [3](#)
45. Park, J., Xu, C., Yang, S., Keutzer, K., Kitani, K.M., Tomizuka, M., Zhan, W.: Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. In: TInternational Conference on Learning Representations (2022) [10](#)
46. Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3967–3976 (2019) [3](#)
47. Phillion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: European Conference on Computer Vision. pp. 194–210 (2020) [3](#), [4](#)

48. Qin, Z., Li, X.: Monoground: Detecting monocular 3d objects from the ground. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3793–3802 (2022) [3](#)
49. Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distribution network for monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8555–8564 (2021) [3](#)
50. Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. The British Machine Vision Conference (2018) [3](#)
51. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9627–9636 (2019) [3](#)
52. Tian, Z., Chen, P., Lai, X., Jiang, L., Liu, S., Zhao, H., Yu, B., Yang, M.C., Jia, J.: Adaptive perspective distillation for semantic segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(2), 1372–1387 (2022) [3](#)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems (2017) [3](#)
54. Wang, T., Xinge, Z., Pang, J., Lin, D.: Probabilistic and geometric depth: Detecting objects in perspective. In: Conference on Robot Learning. pp. 1475–1485 (2022) [3](#)
55. 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) [1](#), [3](#)
56. 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 (2022) [1](#), [3](#), [10](#)
57. Wang, Y., Solomon, J.M.: Object dgcnn: 3d object detection using dynamic graphs. Advances in Neural Information Processing Systems pp. 20745–20758 (2021) [1](#)
58. Wang, Y., Zhou, W., Jiang, T., Bai, X., Xu, Y.: Intra-class feature variation distillation for semantic segmentation. In: European Conference on Computer Vision. pp. 346–362. Springer (2020) [3](#)
59. 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](#), [11](#)
60. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10687–10698 (2020) [3](#)
61. 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) [3](#), [10](#)
62. Yang, Z., Li, Z., Jiang, X., Gong, Y., Yuan, Z., Zhao, D., Yuan, C.: Focal and global knowledge distillation for detectors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4643–4652 (2022) [2](#), [3](#)
63. 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](#), [9](#)
64. Yuan, L., Tay, F.E., Li, G., Wang, T., Feng, J.: Revisiting knowledge distillation via label smoothing regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3903–3911 (2020) [2](#), [3](#)

65. Zeng, J., Chen, L., Deng, H., Lu, L., Yan, J., Qiao, Y., Li, H.: Distilling focal knowledge from imperfect expert for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 992–1001 (2023) [3](#)
66. Zhang, H., Yang, D., Yurtsever, E., Redmill, K.A., Özgüner, Ü.: Faraway-frustum: Dealing with lidar sparsity for 3d object detection using fusion. In: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC). pp. 2646–2652 (2021) [5](#)
67. Zhang, L., Dong, R., Tai, H.S., Ma, K.: Pointdistiller: Structured knowledge distillation towards efficient and compact 3d detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 21791–21801 (2023) [3](#)
68. Zhang, P., Kang, Z., Yang, T., Zhang, X., Zheng, N., Sun, J.: Lgd: label-guided self-distillation for object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 3309–3317 (2022) [2](#), [4](#), [6](#)
69. Zhang, Y., Dong, Z., Yang, H., Lu, M., Tseng, C.C., Du, Y., Keutzer, K., Du, L., Zhang, S.: Qd-bev: Quantization-aware view-guided distillation for multi-view 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3825–3835 (2023) [3](#)
70. Zhao, H., Zhang, Q., Zhao, S., Zhang, J., Tao, D.: Bevsimdet: Simulated multi-modal distillation in bird’s-eye view for multi-view 3d object detection. arXiv preprint arXiv:2303.16818 (2023) [11](#)
71. 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) [4](#), [11](#)
72. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019) [3](#)
73. 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) [9](#), [10](#)