

RayD3D: Distilling Depth Knowledge Along the Ray for Robust Multi-View 3D Object Detection

Rui Ding¹, Zhaonian Kuang¹, Zongwei Zhou², Meng Yang^{1*}, Xihu Zheng^{3*}, Gang Hua⁴

¹ Xi'an Jiaotong University, Xi'an, China

² Horizon Robotics, Beijing, China

³ Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

⁴ Amazon, Bellevue, United States

{dingrui, znkwong}@stu.xjtu.edu.cn, ZongweiZhou@horizon.ai, mengyang@mail.xjtu.edu.cn, xinhuzheng@hkust-gz.edu.cn, ganghua@gmail.com

Abstract

Multi-view 3D detection with bird's eye view (BEV) is crucial for autonomous driving and robotics, but its robustness in real-world is limited as it struggles to predict accurate depth values. A mainstream solution, cross-modal distillation, transfers depth information from LiDAR to camera models but also unintentionally transfers depth-irrelevant information (e.g. LiDAR density). To mitigate this issue, we propose RayD3D, which transfers crucial depth knowledge along the ray: a line projecting from the camera to true location of an object. It is based on the fundamental imaging principle that predicted location of this object can only vary along this ray, which is finally determined by predicted depth value. Therefore, distilling along the ray enables more effective depth information transfer. More specifically, we design two ray-based distillation modules. Ray-based Contrastive Distillation (RCD) incorporates contrastive learning into distillation by sampling along the ray to learn how LiDAR accurately locates objects. Ray-based Weighted Distillation (RWD) adaptively adjusts distillation weight based on the ray to minimize the interference of depth-irrelevant information in LiDAR. For validation, we widely apply RayD3D into three representative types of BEV-based models, including BEVDet, BEVDepth4D, and BEVFormer. Our method is trained on clean NuScenes, and tested on both clean NuScenes and RoboBEV with a variety types of data corruptions. Our method significantly improves the robustness of all the three base models in all scenarios without increasing inference costs, and achieves the best when compared to recently released multi-view and distillation models.

Introduction

Multi-view 3D object detection with bird's eye view (BEV) is widely applied in autonomous driving and robotics navigation. They offer a unified feature representation space for multi-view images and perform well on benchmarks with clean data. However, their performance drops significantly in real-world scenarios especially when data corruptions occur such as fog and snow. Figure 1(c) show that a key bot-

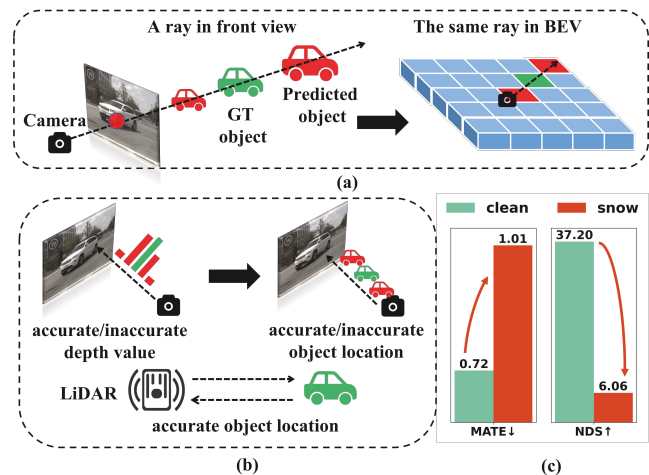


Figure 1: Ray prior for cross-modal distillation. (a) The line projecting from the camera to true location of an object forms a ray in both front view and BEV. Predicted location of this object can only vary along this ray, which is finally determined by predicted depth value. (b) LiDAR has accurate depth. However, inaccurate depth from camera leads to inaccurate location of object along the ray. Therefore, distilling along the ray enables more effective depth transfer. (c) The accuracy (MATE) of predicted depth drops from 0.72 to 1.00 when data corruptions occur, consequently, the accuracy (NDS) of 3D detection drops from 37.20 to 6.06.

tleneck is that they struggle to predict accurate depth values and data corruptions further worsen the problem.

Since LiDAR provides precise depth information (Zhou and Tuzel 2018; Sun et al. 2020), distilling LiDAR features to the camera offers a promising solution to mitigate this problem. However, existing cross-modal distillation methods (Gong et al. 2022; Ding, Yang, and Zheng 2024) simply force camera features to mimic LiDAR features or require a teacher network with the same structure as the student (Huang et al. 2024; Wu et al. 2023). In these methods, LiDAR models unintentionally transfer depth-irrelevant information such as point cloud density and reflection intensity,

*Corresponding author.




Figure 2: Visual examples of our RayD3D. The accuracy of multi-view 3D detection drops significantly in real-world scenarios. Our method consistently demonstrates strong robustness on both clean data and various types of data corruptions, whether data corruptions are seen or not during training.

to camera models. Therefore, these methods fail to effectively transfer crucial depth information to camera models.

In this paper, we apply the ray prior in Figure 1 to more effectively transfer depth information from LiDAR to camera. The ray refers to the line projecting from the camera to true location of an object during imaging. In multi-view 3D detection, predicted location of this object can only vary along this ray in both front view and BEV view as shown in Figure 1(a). The location of this object is finally determined by predicted depth value, as shown in Figure 1(b). Inaccurate depth values lead to inaccurate locations of the object along the ray, thereby, it reduces the detection accuracy. Therefore, distilling along the ray enables more effective depth information transfer from LiDAR to camera model.

The ray prior is a fundamental principle during imaging, which can be used in different ways in relevant 3D tasks such as NeRF (Mildenhall et al. 2021) and pose estimation (Zhan et al. 2022). The recent RayDN (Liu et al. 2024) uses rays in multi-view models to handle redundant predictions. However, they applied the ray in single-modal images.

In this paper, we propose a novel method, namely RayD3D, to more effectively distill depth knowledge along the ray for robust 3D object detection in real-world scenarios. We unify LiDAR and camera features in the BEV space and successfully apply the ray prior into cross-modal distillation. It can be combined with previous ray-based multi-view methods to achieve better performance. Our method includes two ray-based distillation modules. (I) **Ray-based Contrastive Distillation (RCD)** module incorporates contrastive learning into distillation by sampling along the ray to learn how LiDAR accurately locates objects. (II) **Ray-based Weighted Distillation (RWD)** module adaptively adjusts distillation weights based on rays to minimize interference from depth-irrelevant information in LiDAR.

The RCD module incorporates contrastive learning (Sun et al. 2023) by constructing positive and negative sample pairs along the ray. Camera BEV features with accurate depth are sampled as positives, while other BEV features along the ray are negatives. By contrasting them with LiDAR features, the camera model is enhanced to learn LiDAR’s knowledge to distinguish accurate and inaccurate object locations along rays. Contrastive learning ensures that the model focuses on learning depth knowledge from Li-

DAR rather than simply mimicking LiDAR features.

The RWD module adaptively adjusts distillation weights based on rays to control the amount of transferred LiDAR information. For rays with large differences between camera and LiDAR features, we increase weight to transfer more depth information for correction. For rays with small differences, we reduce weight to avoid interfering camera model’s own localization ability. The adaptive weight balances the benefits of transferred depth information and the interference from depth-irrelevant information.

Our RayD3D is compatible with all dominant types of BEV detection models, including LSS-based (Pillion and Fidler 2020), Transformer-based (Li et al. 2022), and temporal models (Huang and Huang 2022). For validation, we widely apply RayD3D into three representative models: BEVDet (Huang et al. 2021), BEVFormer (Li et al. 2022), and BEVDepth4D (Li et al. 2023). Our method is trained on nuScenes with clean data, and evaluated on both nuScenes and RoboBEV with a variety types of data corruptions.

Extensive experiments show that our RayD3D significantly improves the accuracy and robustness of the three base models in all scenarios such as the examples in Figure 2. The performance of our method further improves when trained on corrupted data. Our method also achieves the best performance when compared to recently released multi-view and distillation models with the same setting.

Our main contributions are highlighted as follows.

- We are the first to apply the ray prior for cross-modal knowledge distillation, which enables more effective depth information transfer from LiDAR to camera.
- We propose two novel ray-based distillation modules to enhance the robustness of 3D detection. Our method is compatible with dominant types of BEV-based models.
- Our method significantly boosts the robustness of three representative base models on both clean nuScenes and RoboBEV with a variety types of data corruptions.

Related Work

Multi-View 3D Detection

Multi-View 3D object detection (Wang et al. 2021, 2022) predicts objects’ dimensions, locations, and orientations

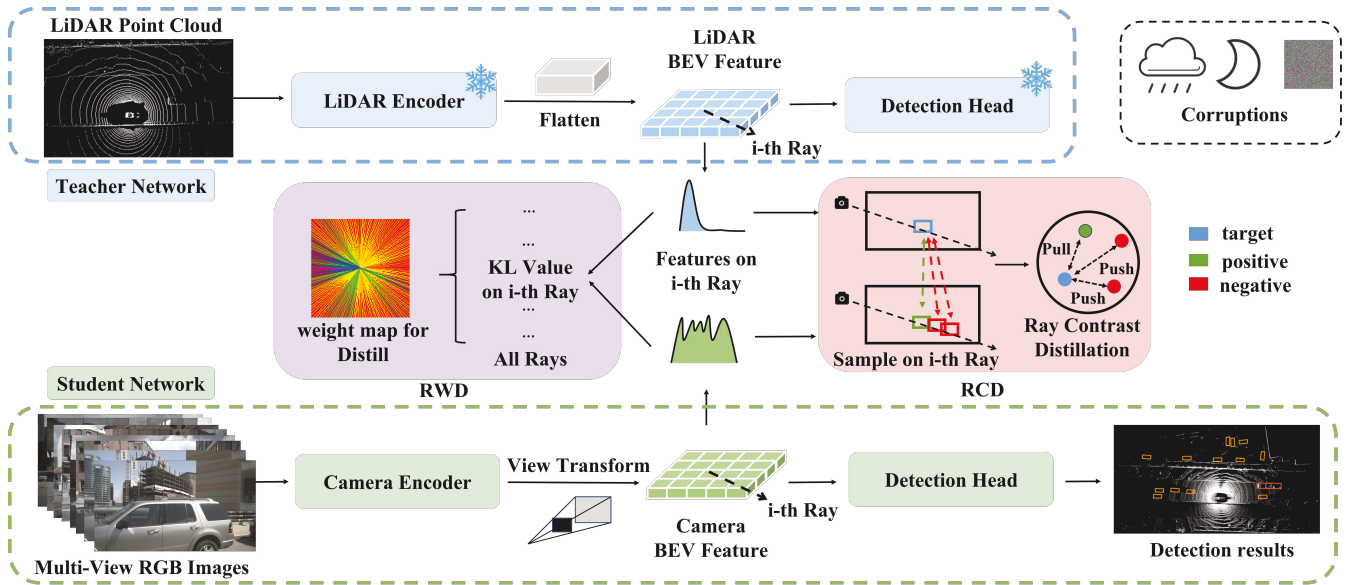


Figure 3: Overview of our RayD3D framework. It includes a teacher network, a student network, and two novel ray-based distillation modules. First, we train the LiDAR-based teacher network and freeze its parameters. The student network is then initialized randomly and trained with our **RCD** and **RWD** distillation modules. During inference, only the student network is used and evaluated on both clean and corrupted data. It significantly enhances the accuracy and robustness in real-world scenarios without increasing inference costs.

from images. Recently, BEV-based models have become dominant and the core is transforming images into BEV. Thus, these models can be divided into two categories: (I) LSS-based (Lift, Splat, Shoot) (Phillion and Fidler 2020; Li et al. 2023). They lift camera features to 3D space and project them onto the BEV. However, inaccurate depth estimation affects detection accuracy. (II) Transformer-Based (Li et al. 2022). They use deformable attention to dynamically sample features from images and map them to the BEV. This allows a single pixel to map to multiple BEV locations, which can also result in inaccurate locations. In addition, models like BEVDepth4D (Huang and Huang 2022), incorporate temporal information and achieve detection accuracy close to LiDAR-based detectors.

Data Corruptions

However, most existing multi-view models cannot well handle data corruptions (Xie et al. 2023; Kong et al. 2024) in real world, such as fog. To enhance robustness, researchers use denoising methods (Guo et al. 2020) or simulate severe weather (Qian et al. 2024; Gupta et al. 2024). These strategies help models perform well on specific data corruptions, however, may fail on other types of data corruptions in real world. It is impractical to denoise or simulate all types of unexpected data corruptions in real world.

Cross-Modal Knowledge Distillation

Knowledge distillation was originally proposed for model compression (Hinton 2015) and then widely applied in various domains (Yang et al. 2022; Chen et al. 2017). In multi-view 3d detection tasks, cross-modal knowledge distilla-

tion methods (Klingner et al. 2023; Zhao et al. 2024) like BEVDistill transfer information from LiDAR teacher network to camera student network. LiDAR-based teacher network can compensate for the limitation of camera model. To reduce modality gap, recent methods (Kim et al. 2024; Zhou et al. 2023; Huang et al. 2024; Wu et al. 2023) transform features onto the BEV and train student BEV features to mimic those of the teacher. However, these methods fail to effectively transfer crucial depth information and do not fully consider camera model robustness in real world. In contrast, we successfully apply the ray prior into cross-modal distillation to improve the robustness of camera models.

Method

Overall Framework

Figure 3 shows our framework, which is designed to improve performance of the multi-view models in all scenarios. It includes three main components: (I) The teacher network takes LiDAR input to generate BEV features; (II) the student network takes N_{camera} camera images as input and transforms features into BEV perspective; (III) two novel ray-based distillation modules, Ray-based Contrastive (RCD) and Ray-based Weighted (RWD) transfer crucial depth information to enhance the camera model’s performance.

More specifically, during distillation process, we use two convolutional layers to resize feature maps of camera to the same size as LiDAR feature maps. BEV feature maps of the camera $C \in R^{D \times H \times W}$ and BEV feature maps of the LiDAR $L \in R^{D \times H \times W}$ are used for distillation, where D , H , and W represent the dimensions of the BEV feature maps.

For distilling along the rays, we evenly divide the BEV feature map into N_{ray} rays.

Ray-Based Contrastive Distillation (RCD)

The RCD module guides the model to learn how LiDAR distinguishes accurate and inaccurate locations along rays. By leveraging the ray prior during the front view to BEV transformation (Zhang et al. 2023; Liu et al. 2024; Chu et al. 2024), it transfers depth information more effectively than directly mimicking LiDAR features. This is achieved by sampling positive and negative pairs along the rays, as shown in Figure 3.

For positive samples on the i -th ray, we sample features from the object’s foreground region, with the camera and LiDAR features at the same j -th position denoted as $C_{i,j}$ and $L_{i,j}$, respectively. It ensures that the sampled positive features correspond to the accurate object locations.

For negative samples, we prioritize sampling from the inaccurate locations near positive samples, ensuring that they are similar in camera features but differ significantly in depth location. This enhances the effectiveness of distinguishing accurate and inaccurate locations. Specifically, We adopt a Gaussian-based sampling strategy to control the probability of negative samples, making the position closer to the positive sample more likely to be selected.

The negative sample from the k -th position on the i -th ray is denoted as $C_{i,k}$. The sampling probability $P(C_{i,k})$:

$$P(C_{i,k}) = \frac{\exp\left(-\frac{d(i,j,k)^2}{2\sigma^2}\right)}{\sum_l \exp\left(-\frac{d(i,j,l)^2}{2\sigma^2}\right)}, \quad (1)$$

where $d(i, j, k)$ represents the distance between samples $C_{i,j}$ and $C_{i,k}$. Smaller σ values increase the sampling probability of position closer to the positive sample. Each ray samples one positive sample and N_{neg} negative samples. Notably, instead of selecting individual pixel features from the BEV feature map, we sample from an $m \times m$ region around each location to capture local context. This helps improve the stability of the loss during distillation.

After constructing the positive and negative pairs, we compute the ray-based contrastive distillation loss L_{RCDs} on the student network and L_{RCDT} on the teacher network, respectively. They are defined as follows:

$$L_{RCDs} = \sum_{i=1}^{N_{ray}} \frac{\exp(L_{i,j} \cdot C_{i,j}/\tau)}{\sum_{k=1}^{N_{neg}+1} \exp(L_{i,j} \cdot C_{i,k}/\tau) + \xi}, \quad (2)$$

$$L_{RCDT} = \sum_{i=1}^{N_{ray}} \frac{\exp(C_{i,j} \cdot L_{i,j}/\tau)}{\sum_{k=1}^{N_{neg}+1} \exp(C_{i,j} \cdot L_{i,k}/\tau) + \xi}, \quad (3)$$

where τ is the temperature parameter to adjust the contrastive strength between positive and negative samples. The factor ξ is set to a constant value $1e-6$ to prevent zero in the denominator. $L_{i,k}$ refers to the negative sample from the k -th position on the i -th ray of LiDAR feature maps. Computing contrastive loss (He et al. 2020; Tian, Krishnan, and Isola 2019) in both the student and teacher networks increases the

diversity of negative examples, further improving the effectiveness of contrastive learning. Finally, the total RCD loss is defined as follows:

$$L_{RCD} = -\log\left(\frac{1}{N_{ray}}(L_{RCDs} + L_{RCDT})\right). \quad (4)$$

As shown in Figure 3, the loss pulls positive pair features closer while pushing negative pair features further apart, improving the model’s ability to distinguish accurate and inaccurate locations. This finally enhances the robustness.

Ray-Based Weighted Distillation (RWD)

The RWD module minimizes interference from depth-irrelevant information by adaptively adjusting distillation weights. The weights are set based on depth estimation accuracy of different regions. However, models like BEVFormer do not explicitly predict depth, so depth error cannot be directly computed. As shown in Figure 3, LiDAR has precise depth information, with high activation of features only at foreground locations along a ray. Camera has inaccurate depth estimation, causing high activation in multiple areas. If the feature distribution of the camera is similar to LiDAR along the ray, it means the foreground object’s location in the camera model is precise.

By analyzing feature differences along each ray, we can determine the accuracy of depth estimation and the amount of depth information needed by camera. To measure the difference between Camera and LiDAR features in each ray, we use Kullback-Leibler (KL) divergence. However, directly calculating KL divergence on the original feature maps is challenging. Feature maps with too many channels make the distillation less stable and difficult to optimize. Therefore, we first generate spatial attention maps for camera and LiDAR, denoted as SC and SL , respectively:

$$SC = \text{softmax}\left(\frac{1}{D} \sum_{d=1}^D |C^d|\right), \quad (5)$$

$$SL = \text{softmax}\left(\frac{1}{D} \sum_{d=1}^D |L^d|\right), \quad (6)$$

where C^d and L^d are the d -th channel of camera and LiDAR features, respectively. $SC \in R^{H \times W}$ and $SL \in R^{H \times W}$ are the spatial attention maps for camera and LiDAR, respectively. This strategy reduces the impact of the modality gap between camera and LiDAR.

Next, we compute feature difference along the i -th ray: $A_i = KL(SL_i, SC_i)$. SL_i and SC_i represent the LiDAR and camera features along the i -th ray. As shown in Figure 4, a large A_i indicates the camera model predicts depth inaccurately, so we increase the distillation weight to transfer more LiDAR depth information for correction. A small A_i means the camera model predicts depth accurately on its own. We reduce the weight to limit interference from depth-irrelevant information. Then, we apply each ray’s weight A_i to all the points along its ray in BEV as shown in Figure 3. The generated weight map ω is as follows:

$$\omega[h, w] = \frac{1}{I((h, w) \in R_i)} \sum_{i=1}^{N_{ray}} A_i \cdot I((h, w) \in R_i), \quad (7)$$




Figure 4: Distillation weight in RWD. For rays with large differences between camera and LiDAR features, we increase weight to transfer more depth information for correction. For rays with small differences, we reduce weight to avoid the interference of depth-irrelevant information on the camera model itself.

where $\omega \in R^{H \times W}$ represents final weight map, and R_i denotes all coordinates along the i -th ray. The indicator function $I((h, w) \in R_i)$ determines whether coordinate $[h, w]$ lies on ray R_i . Inspired by previous works (Wang et al. 2023; Chong et al. 2022), we scale the background weights by a constant factor s to reduce its influence in distillation.

To ensure consistent weighting within each foreground object, all regions inside the same object are assigned the same weight. Specifically, this weight is determined by taking the maximum value among all rays passing through the object. Finally, the ray-based weighted distillation loss L_{RWD} is as follows:

$$L_{RWD} = \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W (\omega[h, w] \cdot |L[h, w] - C[h, w]|). \quad (8)$$

In total, the adaptive weighting strategy enables the camera model to learn reliable depth information from LiDAR while minimizing interference. It significantly enhances depth estimation ability of the camera model, reducing the impact of corruptions on detection accuracy. In addition, it avoids overfitting to LiDAR features.

Total Loss

In addition to RCD and RWD, we also incorporate the classic response distillation loss (Ding, Yang, and Zheng 2024). Response distillation allows the student to be close to the teacher’s output. The loss L_{RES} is as follows:

$$L_{RES} = \frac{1}{N_{heads}} \sum_{n=1}^{N_{heads}} (Y_{L,n} - Y_{C,n}), \quad (9)$$

where N_{heads} denotes the number of prediction heads, and $Y_{L,n}$ and $Y_{C,n}$ represent the features of the n -th head for LiDAR and camera, respectively. We also use the original loss L_{SRC} for 3D detection in the student network. Our total loss function is as follows:

$$L = L_{RCD} + L_{RWD} + L_{SRC} + L_{RES}. \quad (10)$$

This loss can be applied to all types of BEV models in the literature, enhancing their robustness in real-world scenarios without increasing inference costs.

Experiments

Setting

Dataset and Metrics Our method is evaluated on the large-scale autonomous driving benchmarks nuScenes (Caesar et al. 2020). nuScenes includes clean images captured with six camera views and a 32-beam LiDAR, which consists of 700 training, 150 validation, and 150 testing scenes. We use the nuScenes Detection Score (NDS) and mean Average Precision (mAP) as metrics.

RoboBEV (Xie et al. 2023) includes eight types of data corruptions based on nuScenes including Brightness, low light, Fog, Snow, Motion Blur, Color Quant, Camera Crash, and Frame Lost. The mean Resilience Rate (**mRR**) is employed to evaluate the average performance of N types of data corruptions as $mRR = \frac{1}{N} (NDS_i / NDS_{clean})$, where NDS_{clean} and NDS_i refer to NDS values on clean data and corrupted data, respectively. Notably, RoboBEV only adds corruptions to nuScenes validation, so there is no risk of information leakage.

Network Configurations For validation, we select three representative base models BEVDet (Huang et al. 2021), BEVFormer (Li et al. 2022), and BEVDepth4D (Li et al. 2023). For the teacher network, we choose the widely used CenterPoint (Yin, Zhou, and Krahenbuhl 2021) for fair comparison. Both teacher and student models transform inputs into BEV feature maps for distillation and detection.

The parameters for the teacher and student networks follow those of the original models unless otherwise specified. All base models use ResNet-50 as backbone. For BEVDet and BEVDepth4D, the image size is downsampled to 256×704, while for BEVFormer, it is set to 800×450. The BEV grid size of BEVFormer is 50×50. BEVDepth4D-r101 use backbone ResNet-101 with image size 512*1408.

The KL temperature is 0.07, and the number of negatives is 5. We fix the random seed to 0, and the variance across multiple runs is below 0.2 NDS. With our distillation modules, the training time slightly increases when training for 24 epochs on 8×3090 GPUs, while the inference time not increase.

Evaluation

We train all models on clean NuScenes and test them on both NuScenes and RoboBEV. It aims to verify the accuracy on clean data, and the robustness on corrupted data in real world even though data corruptions are not seen during training. We then compare our method with recent multi-view detection and cross-modal distillation methods. In all experiments, clean LiDAR data are used to ensure accurate depth information transfer. Since both teacher and student are trained on clean nuScenes training set and evaluated on corrupted RoboBEV validation set, there is no risk of information leakage. In addition, we also show the results that trained and tested models with data corruptions, which aims to further verify the robustness of our method on corrupted data.

Comparison to Base Models Table 1 shows that our method significantly improves the accuracy of the three base

Method	percentage (NDS_i/NDS_{clean}) % (RoboBEV)									NDS/mAP (nuScenes)	
	Brightness	Camera Crash	Color Quant	Fog	Frame Lost	Low Light	Motion Blur	Snow	mRR	Val	Test
BEVDet	67.9	66.3	65.4	64.9	52.5	29.2	55.1	16.3	52.2	37.4/29.6	38.4/28.9
+ours	69.6	66.3	66.0	66.4	57.0	29.5	59.2	23.8	54.7	41.9/32.7	43.1/33.7
BEVDepth4D	73.3	68.2	71.3	76.6	71.0	41.8	64.8	22.4	61.1	48.4/36.4	49.0/38.2
+ours	76.7	68.6	72.8	79.0	73.3	49.3	66.7	30.4	64.6	50.7/39.3	52.0/40.3
BEVFormer	89.5	67.8	76.9	86.2	53.4	27.9	61.6	33.5	62.1	35.4/25.2	-
+ours	91.9	67.9	78.5	87.0	59.2	32.2	59.8	42.0	64.8	41.4/31.5	-
BEVDepth4D-r101	67.8	69.9	61.4	79.9	73.5	19.3	57.1	18.5	55.9	53.5/41.2	55.0/43.3
+ours	72.1	70.5	62.4	80.0	71.6	27.3	61.1	22.1	58.4	55.3/45.4	57.6/47.3

Table 1: Comparison to base models on nuScenes and RoboBEV. Our method significantly improves the accuracy (**NDS/mAP**) of base models on clean data and enhances the robustness (**mRR**) of all eight types of data corruptions. We use the same setting FP16 with official codes of BEVDet and BEVDepth4D to ensure a fair comparison with recent cross-modal distillation methods DistillBEV (Wang et al. 2023) and VexKD (Yuzhe et al. 2025).

Method	Type	Teacher	Student	Backbone	Frame	NDS/mAP	mRR
PETR* (Liu et al. 2022)	Multi-View	-	-	ResNet-50	1	36.7/31.7	61.3
BEVerse* (Zhang et al. 2022)	Multi-View	-	-	SwinT	1+3	46.6/32.1	48.6
Sparse4D* (Lin et al. 2022)	Multi-View	-	-	ResNet-101	1+3	54.4/44.0	55.0
SoloFusion* (Park et al. 2023)	Multi-View	-	-	ResNet-50	1+16	48.5/38.7	64.4
DualBEV (Li et al. 2024)	Multi-View	-	-	ResNet-50	1+1	50.4/38.0	-
Unidistill (Zhou et al. 2023)	Distill	L+C	BEVDet	ResNet-50	1	39.3/29.6	50.6
DistillBEV (Wang et al. 2023)	Distill	L	BEVDet	ResNet-50	1	40.6/31.4	52.4
VexKD (Yuzhe et al. 2025)	Distill	L+C	BEVDet	ResNet-50	1	40.6/ 34.7	-
Ours	Distill	L	BEVDet	ResNet-50	1	41.9/32.7	54.7
TiG-BEV (Huang et al. 2022)	Distill	L	BEVDepth4D	ResNet-50	1+1	46.1/36.6	62.2
DistillBEV (Wang et al. 2023)	Distill	L	BEVDepth4D	ResNet-50	1+1	50.3/38.3	62.6
LabelDistill (Kim et al. 2024)	Distill	L	BEVDepth4D	ResNet-50	1+1	50.1/37.9	-
Ours	Distill	L	BEVDepth4D	ResNet-50	1+1	50.7/39.3	64.6
DistillBEV (Wang et al. 2023)	Distill	L	BEVDepth4D	ResNet-101	1+1	53.1/42.7	58.4
Ours	Distill	L	BEVDepth4D	ResNet-101	1+1	54.8/44.0	59.4

Table 2: Comparison to recent models. ‘L+C’ and ‘L’ indicate that the teacher network uses a fusion model (LiDAR + Camera) and a single LiDAR model, respectively. * refers to the results are obtained from RoboBEV official github (Xie et al. 2023). ‘1+1’ indicates using current and one previous frame. For fair comparison, we use the version of LabelDistill with 2 frames.

models on clean data and all types of data corruptions. Although BEVFormer does not explicitly predict depth, it still struggles with inaccurate localization when transforming the front view to BEV. Our method well boosts its performance in all scenarios. The reason behind is that precise depth helps maintain stable BEV projections even when corruptions occur, leading to more robust result.

Comparison to Recent Multi-View Models Table 2 shows that our method achieves higher accuracy on both clean data (NDS/mAP in nuScenes) and data corruptions (mRR in RoboBEV) when compared to recent multi-view models, even though some multi-view detection models use stronger backbones and higher image resolutions.

Comparison to Recent Distillation Models Table 2 shows that with the same base model, our method outperforms all recent cross-modal distillation methods on both clean and corrupted data. Notably Unidistill and VexKD use fusion models as teachers, while our method uses only Li-

DAR models. Our model still achieves better performance. This is because these cross-modal methods simply force camera features to mimic LiDAR features. In contrast, our RCD and RWD modules effectively guide the camera to learn relevant LiDAR depth information.

Train and Test with Data Corruptions Our method is robust to various data corruptions in real-world scenarios, whether they are seen or not during training. Table 1 has verified the robustness when training on clean data and test on corrupted data. To further validate, we simulate the eight types of data corruptions during training, and test on RoboBEV with data corruptions in Table 3. In this case, the impact of data corruptions on depth is known during training, allowing our model to more effectively learn from clean LiDAR information. Our method improves the robustness of the base model BEVDet by 5.1%, and also outperforms the representative distillation method DistillBEV.

Method	percentage (NDS_i/NDS_{clean}) % (RoboBEV)									NDS/mAP (nuScenes)
	Brightness	Camera Crash	Color Quant	Fog	Frame Lost	Low Light	Motion Blur	Snow	mRR	Clean
BEVDet	91.7	69.2	91.6	93.7	53.0	55.2	71.1	81.2	75.8	38.6/29.9
+DistillBEV	93.6	70.0	92.6	94.6	58.4	53.6	72.4	85.1	77.5	40.9/31.7
+ours	96.3	70.0	94.8	95.8	60.1	61.8	78.0	90.7	80.9	41.6/32.0
BevDepth4D	93.5	72.3	92.6	95.0	85.4	64.2	78.9	83.6	83.2	49.1/37.1
+DistillBEV	94.8	71.6	94.4	96.2	85.8	67.3	79.1	88.0	84.7	50.0/38.6
+ours	96.1	72.1	95.6	96.8	85.8	69.4	81.8	92.6	86.3	50.8/38.9

Table 3: Train and test with data corruptions. The results further verify the robustness of our method, either data corruptions are seen or not during training.

	RCD	RWD	NDS/mAP	mRR
a			37.4/29.6	52.2
b	✓		40.6/31.8	53.7
c		✓	41.7/32.0	54.1
d	✓	✓	41.9/32.7	54.7

Table 4: Ablation study on our two distillation modules.

Sample negative	Sample in teacher	NDS/mAP	mRR
Random	✓	39.4/29.9	53.2
		40.0/30.1	53.0
Gaussian (ours)	✓	40.6/31.8	53.0
		40.6/31.8	53.7

Table 5: Ablation study on sampling strategies in RCD.

Ablation Studies

We conduct ablation studies to verify the effectiveness of each module. All experiments use BEVDet as base model with an image size of 256×704 and a ResNet-50 backbone.

Two Ray-Based Distillation Modules Table 4 validates the effectiveness of our two distillation modules RCD and RWD. The results indicate that each module independently enhances the robustness of the base model, and their combination achieves much higher performance. The RCD module helps the model learn how LiDAR distinguishes accurate and inaccurate locations, while the RWD module minimizes interference from depth-irrelevant information. Notably, two modules can be well applied to all mainstream types of BEV-based models to enhance their robustness.

Sampling Strategies in RCD In the RCD module, we sample negative samples for contrastive learning. We test two sampling strategies: (1) "Random" random samples in background regions along the ray, and (2) "Gaussian" uses a Gaussian distribution to increase sampling probability of regions closer to positive samples.

Table 5 shows that the "Gaussian" strategy achieves better accuracy and robustness, which helps the network more effectively to distinguish nearby positive and negative samples. The results also show that sampling in the teacher network further improves the robustness, because it increases the variety of negative samples.

Strategies	NDS/mAP	mRR
Cosine similarity	40.2/31.1	53.2
JS divergence	41.0/31.8	53.4
KL divergence (ours)	41.7/32.0	54.1

Table 6: Ablation studies on calculation strategies of cross-modal feature difference in RWD.

Calculation of Feature Difference in RWD In our weight strategy, we calculate the feature distribution difference between the camera and LiDAR along each ray. We test several strategies. "Cosine similarity" is simple but cannot capture high-dimensional feature differences well. "JS divergence" is less sensitive to small variations between nearby positive and negative samples. Finally, we chose "KL divergence", because it well balances computational cost and performance. Table 6 shows that "KL divergence" achieves the highest accuracy and robustness.

Visualization

Figure 2 visualizes the detection results of both the representative base model BEVDet and our method on clean data and various types of data corruptions. The results demonstrate that our method consistently detects more objects under all corruption types. However, severe corruption causes information to completely lose in a few regions, leading to lower accuracy even with distillation.

Conclusion

In this work, we verify that the robustness of multi-view 3D object detection models is limited by inaccurate depth prediction, especially when data corruptions occur in real-world. We first apply the ray prior for cross-modal knowledge distillation and propose a cross-modal distillation method with two novel ray-based modules. It focuses on crucial depth information while mitigating depth-irrelevant information from LiDAR to camera models. Experiments verify the accuracy and robustness of our method on both NuScenes with clean data and RoboBEV with data corruptions. Our method is compatible with all mainstream types of BEV detection models and has the potential to achieve higher accuracy with stronger base models in the future. We will further extend this work to cross-city settings, broader datasets, and more sensor types.

Acknowledgements

This work was supported by the National Science Foundation of China (NSFC) under Grant 62088102, Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China under Grant JYB2025XDXM504, NSFC under Grants 62373298 and U24A20252, and Nansha Key Science and Technology Project under Grant 2023ZD006.

References

- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.
- Chen, G.; Choi, W.; Yu, X.; Han, T.; and Chandraker, M. 2017. Learning efficient object detection models with knowledge distillation. *Advances in neural information processing systems*, 30.
- Chong, Z.; Ma, X.; Zhang, H.; Yue, Y.; Li, H.; Wang, Z.; and Ouyang, W. 2022. MonoDistill: Learning Spatial Features for Monocular 3D Object Detection. In *International Conference on Learning Representations*.
- Chu, X.; Deng, J.; You, G.; Duan, Y.; Li, Y.; and Zhang, Y. 2024. RayFormer: Improving Query-Based Multi-Camera 3D Object Detection via Ray-Centric Strategies. *arXiv preprint arXiv:2407.14923*.
- Ding, R.; Yang, M.; and Zheng, N. 2024. Selective Transfer Learning of Cross-Modality Distillation for Monocular 3D Object Detection. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Gong, Y.; Khurana, S.; Rouditchenko, A.; and Glass, J. 2022. Cmkd: Cnn/transformer-based cross-model knowledge distillation for audio classification. *arXiv preprint arXiv:2203.06760*.
- Guo, C.; Li, C.; Guo, J.; Loy, C. C.; Hou, J.; Kwong, S.; and Cong, R. 2020. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1780–1789.
- Gupta, H.; Kotlyar, O.; Andreasson, H.; and Lilienthal, A. J. 2024. Robust Object Detection in Challenging Weather Conditions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 7523–7532.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Hinton, G. 2015. Distilling the Knowledge in a Neural Network. *arXiv preprint arXiv:1503.02531*.
- Huang, J.; and Huang, G. 2022. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*.
- Huang, J.; Huang, G.; Zhu, Z.; Ye, Y.; and Du, D. 2021. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*.
- Huang, L.; Li, Z.; Sima, C.; Wang, W.; Wang, J.; Qiao, Y.; and Li, H. 2024. Leveraging vision-centric multi-modal expertise for 3d object detection. *Advances in Neural Information Processing Systems*, 36.
- Huang, P.; Liu, L.; Zhang, R.; Zhang, S.; Xu, X.; Wang, B.; and Liu, G. 2022. Tig-bev: Multi-view bev 3d object detection via target inner-geometry learning. *arXiv preprint arXiv:2212.13979*.
- Kim, S.; Kim, Y.; Hwang, S.; Jeong, H.; and Kum, D. 2024. LabelDistill: Label-guided Cross-modal Knowledge Distillation for Camera-based 3D Object Detection. *arXiv preprint arXiv:2407.10164*.
- Klingner, M.; Borse, S.; Kumar, V. R.; Rezaei, B.; Narayanan, V.; Yogamani, S.; and Porikli, F. 2023. 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*, 13343–13353.
- Kong, L.; Xie, S.; Hu, H.; Ng, L. X.; Cottureau, B.; and Ooi, W. T. 2024. Robodepth: Robust out-of-distribution depth estimation under corruptions. *Advances in Neural Information Processing Systems*, 36.
- Li, P.; Shen, W.; Huang, Q.; and Cui, D. 2024. Dualbev: Unifying dual view transformation with probabilistic correspondences. In *European Conference on Computer Vision*, 286–302. Springer.
- Li, Y.; Ge, Z.; Yu, G.; Yang, J.; Wang, Z.; Shi, Y.; Sun, J.; and Li, Z. 2023. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2, 1477–1485.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 1–18. Springer.
- Lin, X.; Lin, T.; Pei, Z.; Huang, L.; and Su, Z. 2022. Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion. *arXiv preprint arXiv:2211.10581*.
- Liu, F.; Huang, T.; Zhang, Q.; Yao, H.; Zhang, C.; Wan, F.; Ye, Q.; and Zhou, Y. 2024. Ray denoising: Depth-aware hard negative sampling for multi-view 3d object detection. In *18th European Conference on Computer Vision (ECCV)*.
- Liu, Y.; Wang, T.; Zhang, X.; and Sun, J. 2022. Petr: Position embedding transformation for multi-view 3d object detection. In *European Conference on Computer Vision*, 531–548. Springer.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Park, J.; Xu, C.; Yang, S.; Keutzer, K.; Kitani, K. M.; Tomizuka, M.; and Zhan, W. 2023. Time Will Tell: New Outlooks and A Baseline for Temporal Multi-View 3D Object Detection. In *ICLR*.
- Philion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision–ECCV 2020: 16th European*

- Conference, Glasgow, UK, August 23–28, 2020, *Proceedings, Part XIV 16*, 194–210. Springer.
- Qian, C.; Guo, Y.; Mo, Y.; and Li, W. 2024. WeatherDG: LLM-assisted Procedural Weather Generation for Domain-Generalized Semantic Segmentation. *arXiv:2410.12075*.
- Sun, J.; Cao, Y.; Chen, Q. A.; and Mao, Z. M. 2020. Towards robust {LiDAR-based} perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *29th USENIX Security Symposium (USENIX Security 20)*, 877–894.
- Sun, J.; Zheng, H.; Zhang, Q.; Prakash, A.; Mao, Z. M.; and Xiao, C. 2023. Calico: Self-supervised camera-lidar contrastive pre-training for bev perception. *arXiv preprint arXiv:2306.00349*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*.
- Wang, T.; Zhu, X.; Pang, J.; and Lin, D. 2021. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 913–922.
- Wang, Y.; Guizilini, V. C.; Zhang, T.; Wang, Y.; Zhao, H.; and Solomon, J. 2022. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, 180–191. PMLR.
- Wang, Z.; Li, D.; Luo, C.; Xie, C.; and Yang, X. 2023. Distillbev: Boosting multi-camera 3d object detection with cross-modal knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8637–8646.
- Wu, Z.; Wu, Y.; Pu, J.; Li, X.; and Wang, X. 2023. Attention-based depth distillation with 3d-aware positional encoding for monocular 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3, 2892–2900.
- Xie, S.; Kong, L.; Zhang, W.; Ren, J.; Pan, L.; Chen, K.; and Liu, Z. 2023. RoboBEV: Towards Robust Bird’s Eye View Perception under Corruptions. *arXiv preprint arXiv:2304.06719*.
- Yang, C.; Pan, J.; Gao, X.; Jiang, T.; Liu, D.; and Chen, G. 2022. Cross-task knowledge distillation in multi-task recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, 4, 4318–4326.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.
- Yuzhe, J.; Chen, Y.; Yang, L.; Rui, D.; Yang, M.; and Zheng, X. 2025. VeXKD: The Versatile Integration of Cross-Modal Fusion and Knowledge Distillation for 3D Perception. *Advances in Neural Information Processing Systems*, 37: 125608–125634.
- Zhan, Y.; Li, F.; Weng, R.; and Choi, W. 2022. Ray3d: ray-based 3d human pose estimation for monocular absolute 3d localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13116–13125.
- Zhang, H.; Li, H.; Liao, X.; Li, F.; Liu, S.; Ni, L. M.; and Zhang, L. 2023. DA-BEV: Depth aware BEV transformer for 3D object detection. *arXiv preprint arXiv:2302.13002*, 3.
- Zhang, Y.; Zhu, Z.; Zheng, W.; Huang, J.; Huang, G.; Zhou, J.; and Lu, J. 2022. Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving. *arXiv preprint arXiv:2205.09743*.
- Zhao, H.; Zhang, Q.; Zhao, S.; Chen, Z.; Zhang, J.; and Tao, D. 2024. Simdistill: Simulated multi-modal distillation for bev 3d object detection. In *Proceedings of the AAAI conference on artificial intelligence*, 7, 7460–7468.
- Zhou, S.; Liu, W.; Hu, C.; Zhou, S.; and Ma, C. 2023. Uni-Distill: 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*, 5116–5125.
- Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.