Towards Accurate 3D Object Detection in Adverse Weather by Leveraging 4D Radar for LiDAR Geometry Enhancement

Tianxu Tong[1], Xinrun Liu[1], Hongmin Liu[1], Bin Fan[1]*

[1]School of Intelligence Science and Technology, University of Science and Technology Beijing, China

Abstract

3D object detection is a critical component of autonomous driving, yet its performance degrades severely in adverse weather due to the degradation of LiDAR point clouds. While existing LiDAR-4D radar fusion methods enhance robustness by incorporating weather-robust 4D radar data, they often depend on well geometric structures from LiDAR and so struggle to effectively exploit radar data in case of degraded LiDAR data. To tackle this challenge, we propose REL, a novel 4D radar-guided LiDAR geometric enhancement framework. It utilizes 4D radar features to dynamically generate virtual LiDAR points, effectively increasing the density of degraded LiDAR data. Moreover, a Position-Guided Cross Attention (PGCA) module is proposed to enhance the feature representation of virtual points, while an Adaptive Feature Fusion (AFF) module is designed to integrate virtual and real LiDAR features. Extensive experiments on the K-Radar and Vod-Fog datasets demonstrate that REL achieves state-of-the-art 3D object detection performance under diverse adverse weather conditions. Notably, REL improves the overall AP3D by 9.3% on K-Radar and boosts the cyclist class by up to 52.9% 3D mAP under the most severe foggy condition on Vod-Fog.

Code — https://github.com/TongTianxu/REL

Introduction

3D object detection is a fundamental technology for autonomous driving (Song et al. 2024; Gu et al. 2025) and robotic navigation (Naich and Carrión 2024), playing a pivotal role for the safe operation of unmanned platforms in complex environments (Gupta et al. 2024). LiDAR-based methods (Liu et al. 2024b; Zhang et al. 2024b; Ma et al. 2023; Hoang, Bui, and Yoo 2024; Wozniak et al. 2024) have demonstrated promising performance under ideal conditions by leveraging precise 3D point cloud representations. However, the inherent short wavelength characteristics of LiDAR renders it highly vulnerable to adverse weather conditions, leading to significant degradation in data quality (Qi et al. 2024; Oh and Dong-O'Brien 2025), which in turn severely impacts detection accuracy.

*Corresponding author. Email:bin.fan@ieee.org

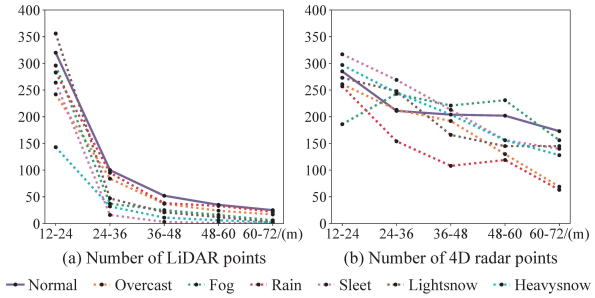

Figure 1: Point distribution of (a) LiDAR and (b) 4D radar in GT boxes under different weather conditions. (a) LiDAR shows a significant reduction in point numbers at long distances and under adverse weather conditions. (b) 4D radar maintains a relatively stable number of points across various weather conditions and distances.

To mitigate the data degradation of LiDAR, detection methods that fuse LiDAR and 4D radar have gained increasing attention in recent years (Wang et al. 2022b; Chae, Kim, and Yoon 2024; Huang et al. 2025b). These approaches enhance detection reliability across various environments by combining LiDAR's high-resolution geometric information with 4D radar's robustness to adverse weather conditions (Kong et al. 2024; Fan et al. 2024). Existing methods predominantly employ hierarchical feature interaction or early fusion strategies, such as 3D-LRF (Chae, Kim, and Yoon 2024) which aggregates 4D radar features within LiDAR voxel neighborhoods, and L4DR (Huang et al. 2025b) that concatenates features from spatially overlapping LiDAR and 4D radar pillars. While these fusion techniques significantly outperform single-modal approaches in 3D object detection, they fail to address the fundamental issue of LiDAR geometric structure degradation. This limitation arises because current fusion methods rely heavily on complete geometric spatial structures provided by LiDAR for the effective extraction of 4D radar information. However, LiDAR sensor experience severe performance degradation under adverse weather conditions (as shown in Figure 1), leading to incomplete geometric structure information. Consequently, 4D radar information cannot be effectively leveraged, resulting in suboptimal fusion performance.

To address these challenges, we propose REL (**R**adar **E**nhanced **L**iDAR), a LiDAR-4D radar fusion detection network designed to leverage 4D radar information for enhancing the geometric structure of degraded LiDAR point clouds, thereby improving 3D object detection under adverse weather. Specifically, we propose Virtual LiDAR Generation (VLG), which generates virtual LiDAR points through weather-robust 4D radar to improve point cloud density and supplement missing geometric structures when LiDAR point clouds degrade. However, due to modality differences between 4D radar and LiDAR, a feature distribution shift exists between the generated virtual and real LiDAR points, which hinders effective multimodal information fusion. Therefore, we design Position-Guided Cross-Attention (PGCA), which uses virtual LiDAR points as queries to aggregate information from neighboring real LiDAR points via cross-attention, thereby enhancing the features of virtual LiDAR points. Finally, Adaptive Feature Fusion (AFF) adaptively fuses virtual and real LiDAR features in voxel space, enabling the model to dynamically focus on the most discriminative information under varying weather conditions. Extensive experiments conducted on the K-Radar (Paek, Kong, and Wijaya 2022) and Vod-Fog (Huang et al. 2025b) datasets demonstrate that REL achieves state-of-the-art performance across diverse weather scenarios.

Our main contributions are as follows:

- We propose REL, a novel LiDAR-4D radar fusion detection network that generates virtual LiDAR points through 4D radar to mitigate LiDAR geometric degradation under adverse weather conditions. The proposed framework is highly flexible and can be easily adapted to various existing fusion methods.
- We design the Position-Guided Cross-Attention (PGCA) module that enhances virtual LiDAR point features through cross-attention mechanisms, effectively addressing the feature distribution discrepancy caused by modality differences between 4D radar and LiDAR.
- Extensive experiments on K-Radar and Vod-Fog benchmarks validate the effectiveness of REL across various adverse weather conditions, achieving state-of-the-art performance on both datasets.

Related Work

LiDAR-based 3D object detection. Methods such as PointPillars and VoxelNet (Lang et al. 2019; Zhou and Tuzel 2018) have exhibited excellent performance under ideal weather conditions but experienced significant performance degradation in adverse weather. To mitigate this issue, various approaches have been proposed. Some studies (Liu et al. 2022; Hahner et al. 2021; Teufel et al. 2022; Scheuble et al. 2024) utilize physical or empirical models to simulate rain and fog noise, aiming to improve the quality of point clouds. Some other methods (Wang et al. 2022c; Roriz et al. 2021; Huang et al. 2023) design filtering mechanisms based on the statistical properties of point clouds to suppress noise. Recently, generative models are employed by (Zhang et al. 2023, 2024d; Matteazzi et al. 2025) to restore degraded point clouds. Additionally, SRKD (Huang et al. 2024) and AWARDistill (Liu et al. 2025) utilize knowledge distillation to transfer knowledge from clear to adverse weather. However, when LiDAR severely degrades under extreme weather conditions, these single-modal methods struggle to fully recover geometric spatial structures, which limits their performance.

Radar-based 3D object detection. Early research mainly adopts 3D radar (Meyer, Kuschk, and Tomforde 2021; Zhang, Nowruzi, and Laganiere 2021; Jiang et al. 2022). Due to the absence of height information, many methods (Li et al. 2022; Yataka et al. 2024) are limited to bird's-eye view (BEV) detection results. Recently, with the advent of 4D radar technology, additional elevation dimension information has become available, significantly enhancing its potential for 3D object detection. RPFA-Net (Xu et al. 2021) designs a self-attention module based on 4D radar pillar features, and SMURF (Liu et al. 2023) proposes multi-scale spatial fusion strategies. As the baseline for K-Radar, RTNH (Paek, Kong, and Wijaya 2022) employs 3D sparse convolution to encode 4D radar tensors. RTNH+ and 4DSRT (Paek, Kong, and Wijaya 2023; Kong, Paek, and Cho 2023) use CA-CFAR (Liu et al. 2024a) and sidelobe filtering techniques to remove noise during 4D radar point cloud generation. Furthermore, (Zhang et al. 2024c) validates radar stability in varying environments through domain shift analysis. Despite these advancements, radar signals inherently suffer from noise interference, achieving precise detection remains challenging when relying solely on 4D radar data.

LiDAR-radar fusion-based 3D object detection. Recent advancements in LiDAR-radar fusion have witnessed innovative architectures for cross-modal feature integration. InterFusion (Wang et al. 2022b) and TransFusion (Zhang et al. 2024a) adopt transformer architectures to enhance cross-modal feature representations. 3D-LRF (Chae, Kim, and Yoon 2024) further introduces a weather-conditional radar-flow gating network to improve robustness under adverse weather. Other notable methods include L4DR (Huang et al. 2025b) which employs a bidirectional feature fusion strategy, and LiRaFusion (Song, Zhao, and Skinner 2024) that proposes a depth-adaptive fusion framework for better feature alignment. Additionally, (Wang et al. 2022a) designs a joint multi-modal and multi-scale fusion strategy to enhance detection accuracy, while RadarDistill (Bang et al. 2024) uses knowledge distillation to transfer information from LiDAR to radar. V2X-R (Huang et al. 2025a) proposes a cooperative LiDAR-4D radar fusion framework based on diffusion models. However, existing methods generally ignore the impact of LiDAR geometric structure degradation under adverse weather on fusion performance. In such cases, the amount of usable information in LiDAR is significantly reduced, which in turn hampers the performance of these fusion-based approaches. Unlike prior works, we propose a 4D radar-guided geometric enhancement method for LiDAR, which dynamically generates virtual LiDAR points to enhance the utilization of multimodal information.

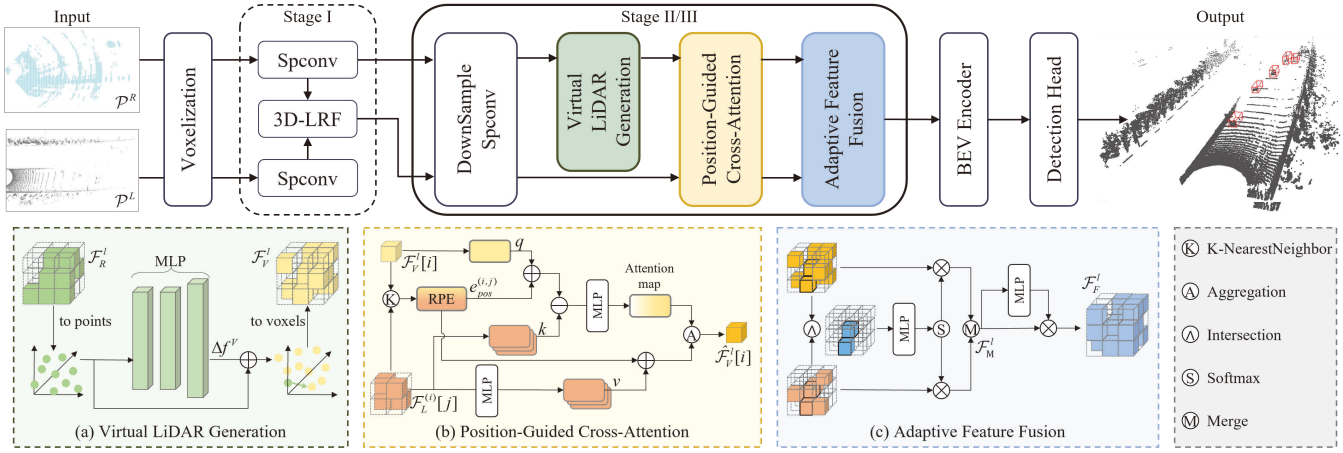

Figure 2: Overall architecture of the proposed REL 3D object detection framework. It consists of three main components: (a) Virtual LiDAR Generation (VLG) generates virtual LiDAR points from 4D radar features, (b) Position-Guided Cross-Attention (PGCA) enhances virtual LiDAR point features by aggregating real LiDAR features through cross-attention and (c) Adaptive Feature Fusion (AFF) module dynamically fuses multi-modal features.

Method

Overall Architecture

Our REL uses 3D-LRF (Chae, Kim, and Yoon 2024) as the baseline for 3D object detection, replacing its fusion modules with our proposed modules. Notably, our method is applicable to existing other fusion methods, which will be demonstrated in the experimental section. The overall architecture of REL is illustrated in Figure 2. The entire network consists of three stages, with our VLG, PGCA, and AFF modules integrated into the second and third stages. These modules are designed to enhance the geometric structure information of LiDAR data under adverse weather conditions by leveraging robust 4D radar features. The rationale behind this design will be thoroughly analyzed in the ablation study. Specifically, REL takes LiDAR point cloud $\mathcal{P}^L$ and 4D radar point cloud $\mathcal{P}^R$ as inputs. After voxelization, high-dimensional feature representations $\mathcal{F}_L^0$ and $\mathcal{F}_R^0$ are obtained. In each stage, LiDAR and 4D radar are initially processed through sparse convolution to obtain $\mathcal{F}_L^l$ and $\mathcal{F}_R^l$ ($l \in \{1, 2, 3\}$). In the first stage, 3D-LRF performs preliminary fusion of 4D radar features with LiDAR, and the fused features $\mathcal{F}_F^1 \in \mathbb{R}^{N_1 \times C_1}$ serve as input to the next LiDAR branch. In the following two stages ($l \in \{2, 3\}$), we employ VLG to generate virtual LiDAR point $\mathcal{P}_V^l \in \mathbb{R}^{M_l \times 3}$ and corresponding features $\mathcal{F}_V^l \in \mathbb{R}^{M_l \times C_l}$ from 4D radar features $\mathcal{F}_R^l$, which supplement the geometric structural information of LiDAR data. Subsequently, the PGCA module enhances the virtual point features, obtaining enhanced features $\hat{\mathcal{F}}_V^l \in \mathbb{R}^{M_l \times C_l}$. The AFF module then performs weighted fusion of $\hat{\mathcal{F}}_V^l$ and $\mathcal{F}_L^l$ in voxel space, generating the final features $\mathcal{F}_F^l \in \mathbb{R}^{N_l \times C_l}$ for the current stage, which is passed to the next LiDAR branch. After each stage, $\mathcal{F}_L^l$ and $\mathcal{F}_F^l$ are encoded into BEV features. The concatenated BEV features from all stages are fed into the detection head, producing the final 3D detection results.

Virtual LiDAR Generation (VLG)

To address the issue of incomplete geometric structures in LiDAR caused by sensor degradation under adverse weather conditions, we propose the Virtual LiDAR Generation (VLG) module to generate virtual points for enhancing the geometric structure of LiDAR by leveraging 4D radar data.

For each non-empty radar voxel, its original spatial coordinates $p_i^R = (x_i^R, y_i^R, z_i^R)$ are obtained through voxel indexing and voxel grid computation, with the corresponding features denoted as $\mathcal{F}_R^l[i]$. VLG takes the 4D radar features as input to predict the offset of virtual LiDAR points by

$$\Delta o_i^V = \Phi_\theta(\mathcal{F}_R^l[i]), \tag{1}$$

where Φ_θ is a three-layer MLP, with the first two layers are followed by ReLU and BatchNorm. VLG has an input dimension of C_l and an output dimension of $3 + C_l$, i.e., $\Phi(\cdot) : \mathbf{R}^{C_l} \rightarrow \mathbf{R}^{3+C_l}$.

The predicted offset Δo_i^V is split into two components: the three-dimensional coordinate offset Δp_i^V and the feature offset Δf_i^V. The generated virtual LiDAR point coordinates and their features can be computed as follows:

$$p_i^V = p_i^R + \Delta p_i^V, \tag{2}$$

$$\mathcal{F}_V^l[i] = \mathcal{F}_R^l[i] + \Delta f_i^V. \tag{3}$$

Through this process, VLG generates virtual LiDAR points with the same quantity as 4D radar points, forming a virtual point set $\mathcal{P}_V^l \in \mathbf{R}^{M_l \times 3}$ and corresponding features $\mathcal{F}_V^l \in \mathbf{R}^{M_l \times C_l}$. These virtual points improve the spatial density and geometric completeness of the LiDAR point cloud, which improves the model's ability to utilize information effectively under adverse weather conditions due to the robustness of 4D radar.

Position-Guided Cross-Attention (PGCA)

Although VLG dynamically generates virtual LiDAR points to supplement missing geometric structures, significant differences between LiDAR and 4D radar result in feature distribution discrepancy between virtual points and real LiDAR points. To address this issue, we propose the Position-Guided Cross-Attention (PGCA) module, which leverages the spatial distribution and contextual information of real LiDAR points to enhance the feature representation of virtual LiDAR points.

For each virtual point $\mathcal{F}_V^l[i]$, we first use KNN to search for K_l ($K_l = \frac{32}{2^{l-1}}$) nearest real LiDAR points $\mathcal{P}_L^{(i)}$. Following Point Transformer (Zhao et al. 2021), we use the Relative Position Encoding (RPE) to model the spatial relationship between virtual point p_i^V and its neighboring real points p_j^L through a linear layer W_p^l:

$$e_{pos}^{(i,j)} = W_p^l(p_i^V - p_j^L). \tag{4}$$

Eq. 4 implicitly captures the spatial relative positions between virtual points and their neighbors, providing geometric constraints for the subsequent attention mechanism.

Subsequently, we construct a cross-attention mechanism that enables virtual points to dynamically aggregate feature information from neighboring real points. Using the virtual LiDAR point features $\mathcal{F}_V^l[i]$ as queries and the features $\mathcal{F}_L^{(i)}[j]$ of the K_l real LiDAR points within the neighborhood as keys, we then incorporate the relative position encoding $e_{pos}^{(i,j)}$ to compute the attention map as:

$$\alpha_{ij} = \textit{softmax}\left(\frac{W_\alpha^l(\mathcal{F}_V^l[i] - \mathcal{F}_L^{(i)}[j] + e_{pos}^{(i,j)})}{\sqrt{d}}\right), \tag{5}$$

where the mapping function W_α^l is a linear layer followed by ReLU. PGCA dynamically suppresses the contribution of low-quality real points through adaptive weighting, ensuring high quality features of virtual points.

Using the attention map, PGCA enhances the feature representation of virtual points by weighted aggregation of neighboring real point features. The enhanced features are computed as follows:

$$\hat{\mathcal{F}}_V^l[i] = \sum_{j=1}^{K_l} \alpha_{ij} \cdot (W_v^l \mathcal{F}_L^{(i)}[j] + e_{pos}^{(i,j)}), \tag{6}$$

where the features $\mathcal{F}_L^{(i)}[j]$ of the K_l real LiDAR points are multiplied by W_v^l to extract value features of real LiDAR features.

The enhanced virtual point features are then mapped to the three-dimensional feature grid $\mathcal{G}^l$ through coordinate transformation $p_i^V \rightarrow \text{Index}(p_i^V)$, which assigns the corresponding $\hat{\mathcal{F}}_V^l[i]$ to the appropriate voxel positions. This ensures that the voxelized representations of virtual and real LiDAR points are spatially aligned, providing a foundation for subsequent multi-modal feature fusion.

Adaptive Feature Fusion (AFF)

To enhance the effectiveness of multi-modal feature fusion between virtual and real LiDAR features, we design an Adaptive Feature Fusion (AFF) module. AFF merges multi-modal features in the aligned voxel space $\mathcal{G}^l$ using a dynamic weighting mechanism.

Firstly, we determine which locations require fusion based on the uniqueness of voxel indices. For each non-empty voxel $i \in \mathcal{I}^l$, if it contains features from only a single source (with unique indices, either virtual LiDAR or real LiDAR), we directly retain the original feature vector. If a voxel i contains both real LiDAR point features $\mathcal{F}_L^l(i)$ and virtual LiDAR point features $\hat{\mathcal{F}}_V^l(i)$ (with repeated indices), we concatenate the corresponding feature vectors and calculate dynamic weights α through a MLP, which is used for fusing cross-modal features as follows:

$$\mathcal{F}_M^l(i) = \alpha \cdot \mathcal{F}_L^l(i) + (1-\alpha) \cdot \hat{\mathcal{F}}_V^l(i). \tag{7}$$

To further enhance feature representation capability, we employ a global attention mechanism to adjust the fused voxel features. First, global average pooling is performed on all voxel features to extract global contextual information. Then, a two-layer MLP (with ReLU in the first layer and Sigmoid in the second layer) is used to learn global attention weights. Finally, the learned weights are applied element-wise to each voxel feature, achieving adaptive feature enhancement based on the global context. The process is formulated as follows:

$$\mathcal{F}_F^l = \mathcal{F}_M^l \odot \left(\text{MLP}\left(\frac{1}{N}\sum \mathcal{F}_M^l\right)\right). \tag{8}$$

After processing through AFF, we obtain the fused features $\mathcal{F}_F^l$ as the output feature representation for the l-th stage, which then serves as the input to the LiDAR branch for the subsequent $(l+1)$-th stage.

BEV Encoder and Detection Head

In the BEV encoder, the LiDAR voxel feature $\mathcal{F}_L^l$ and fusion feature $\mathcal{F}_F^l$ from each stage are encoded to extract BEV features $\mathcal{B}_L^l \in \mathbb{R}^{H\times W\times C}$ and $\mathcal{B}_F^l \in \mathbb{R}^{H\times W\times C}$, respectively. The BEV features from all stages are concatenated to form $\mathcal{B} \in \mathbb{R}^{H\times W\times 6C}$ as input to the detection head.

The detection head consists of a classification head and a regression head. They uses convolutional layers to predict the object class and the 3D bounding box, respectively. The classification head employs focal loss, and the regression head utilizes smooth L_1 loss.

Experiments

Dataset and Evaluation Metrics

K-Radar dataset. K-Radar (Paek, Kong, and Wijaya 2022) is a large-scale autonomous driving perception dataset that covers various weather conditions. It provides multi-modal sensor data, including LiDAR, 4D radar, cameras, and GPS. In this paper, we use only the 64-line LiDAR and 4D radar for training and testing. Following the official split, the training set consists of 17,458 frames, and the test

Methods	Modality	Metric	Total	Normal	Overcast	Fog	Rain	Sleet	Lightsnow	Heavysnow
RTNH (NeurIPS 2022)	4DR	AP_{BEV}	43.2	42.8	58.3	63.7	45.1	28.5	46.5	41.6
		AP_{3D}	15.6	18.0	29.1	18.5	16.6	4.4	25.1	10.2
Pointpillars (CVPR 2019)	L	AP_{BEV}	49.1	48.2	53.0	45.4	44.2	45.9	74.5	53.8
		AP_{3D}	22.4	21.8	28.0	28.2	27.2	22.6	23.2	12.9
RTNH* (NeurIPS 2022)	L	AP_{BEV}	66.3	65.4	87.4	83.8	73.7	48.8	78.5	48.1
		AP_{3D}	37.8	39.8	46.3	59.8	28.2	31.4	50.7	24.6
InterFusion (IROS 2022)	L+4DR	AP_{BEV}	52.9	50.0	59.0	80.3	50.0	22.7	72.2	53.3
		AP_{3D}	17.5	15.3	20.5	47.6	12.9	9.3	56.8	25.7
3D-LRF (CVPR 2024)	L+4DR	AP_{BEV}	73.6	72.3	88.4	86.6	76.6	47.5	79.6	**64.1**
		AP_{3D}	45.2	45.3	55.8	51.8	38.3	23.4	60.2	36.9
REL (Ours)	L+4DR	AP_{BEV}	**75.6**	**74.0**	**88.4**	**90.1**	**84.5**	**52.6**	**84.3**	59.8
		AP_{3D}	**54.5**	**53.0**	**71.5**	**70.3**	**55.4**	**36.6**	**67.4**	**45.7**

Table 1: Quantitative results of LiDAR (L) and 4D radar (4DR)-based 3D object detection methods on K-Radar dataset V1.0. All comparison results are taken from 3D-LRF (Chae, Kim, and Yoon 2024), Best in **bold**, second in underline.

Methods	Metric	Sedan			Bus or Truck		
		Total	Normal	Adverse	Total	Normal	Adverse
InterFusion	AP_{BEV}	68.3	67.5	69.1	50.4	46.5	59.8
3D-LRF		68.4	67.5	69.3	53.2	41.1	67.5
REL		**76.6 (+8.2)**	**75.3 (+7.8)**	**77.8 (+8.5)**	**53.6 (+0.4)**	**41.6 (+0.5)**	**68.7 (+1.2)**
L4DR		68.9	67.7	70.2	51.8	48.2	60.9
REL(with L4DR)		**69.1 (+0.2)**	**67.7 (+0)**	**70.5 (+0.3)**	**58.2 (+6.4)**	**51.7 (+3.5)**	**67.4 (+6.5)**
InterFusion	AP_{3D}	50.9	48.1	52.7	35.5	30.5	44.3
3D-LRF		51.4	50.3	53.0	36.8	29.1	44.7
REL		**53.0 (+1.6)**	**51.2 (+0.9)**	**55.0 (+2.0)**	**40.0 (+3.2)**	**29.8 (+0.7)**	**52.5 (+7.8)**
L4DR		51.2	48.8	54.0	37.6	35.7	43.9
REL(with L4DR)		**52.5 (+1.3)**	**50.5 (+1.7)**	**54.5 (+0.5)**	**41.5 (+3.9)**	**38.3 (+2.6)**	**48.7 (+4.8)**

Table 2: Quantitative results of LiDAR-4D radar fusion-based 3D object detection methods on K-Radar dataset V2.0. All adverse weather conditions are merged into "Adverse". Results include Total, Normal, and Adverse. (+) indicates performance improvement.

set contains 17,536 frames. To assess model performance, we adopt the two standard evaluation versions officially released: V1.0 and V2.0. In version 1.0 (released in 2022), the region of interest (ROI) is defined as [0m, 72m] for the X axis, [-6.4m, 6.4m] for the Y axis, and [-2m, 6m] for the Z axis, with the target category being "Sedan". Version 2.0 (updated in 2023) expands the ROI to [0m, 72m] for the X axis, [-16m, 16m] for the Y axis, and [-2m, 7.6m] for the Z axis, and introduces an additional category, "Bus or Truck". For both versions, performance is evaluated using AP_{BEV} and AP_{3D} metrics with the IoU threshold set at 0.5.

Vod-Fog Dataset. We performed fog simulations on the LiDAR data of VoD dataset (Palffy et al. 2022), introducing five levels of fog density with fog levels $\mathcal{L}$ ranging from 0 to 4 and fog density $\alpha = [0.00, 0.03, 0.06, 0.10, 0.20]$, while keeping the 4D radar data unchanged to simulate various weather conditions, which follows the same data processing approach as L4DR (Huang et al. 2025b). Following the official split, the training and validation sets contain 5,139 and 1,296 frames, respectively. For evaluation, we adopt the official KITTI metrics, testing 3D mAP across three categories: "Car", "Pedestrian", and "Cyclist", under different fog intensities. The results are stratified by detection difficulty levels (easy, moderate, hard) to comprehensively evaluate model performance in foggy scenarios.

Implementation Details

In the PGCA module, the number of neighbors is set hierarchically (with $K = 16$ for the second stage and $K = 8$ for the third stage). For the K-Radar dataset, we train for 11 epochs using batch size 4 and the Adam optimizer, with a learning rate of lr= 1e-3, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. All other parameters are set to the default configurations of 3D-LRF (Chae, Kim, and Yoon 2024). For the Vod-Fog dataset, we train for 100 epochs with batch size 8, and the training strategy follows the original settings of L4DR (Huang et al. 2025b) to ensure a fair comparison. All experiments are conducted on a single RTX 3090 GPU.

Results on K-Radar Dataset

We compare REL with several 3D object detection methods, including PointPillars (Lang et al. 2019), RTNH (Paek, Kong, and Wijaya 2022), InterFusion (Wang et al. 2022b), and 3D-LRF (Chae, Kim, and Yoon 2024), among which InterFusion and 3D-LRF represent state-of-the-art fusion-based approaches. The results in Table 1 highlight the su-

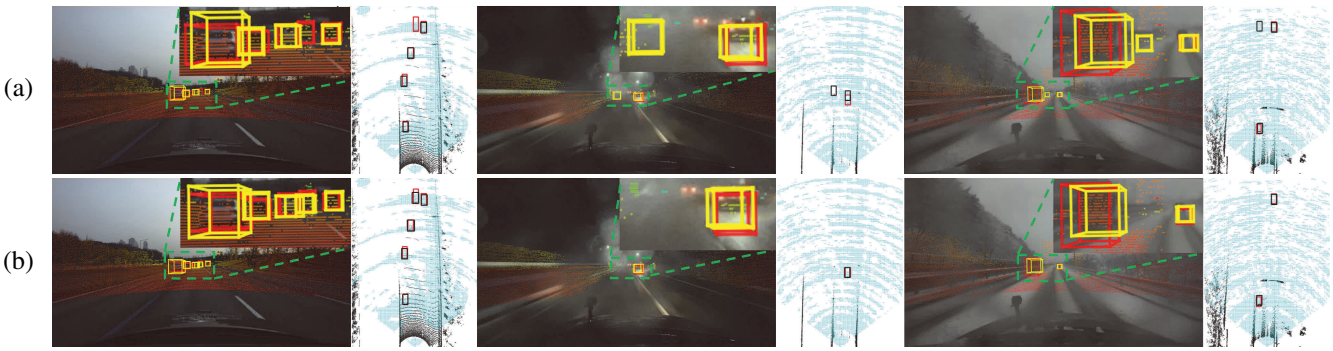

Figure 3: Visualization results of 3D object detection in range view and bird-eye view. The results in the range view show the image and projected LiDAR with red GT boxes and yellow predicted boxes (enlarged results in the upper right corner). The results in bird-eye view show top-view LiDAR and 4D radar heatmap with red GT boxes and black predicted boxes. Each row represents different 3D object detection methods: (a) 3D-LRF, (b) ours. Each column represents different weather conditions: normal, sleet, and heavy snow. It is recommended to enlarge the image for better visualization of the details.

Fog level	Methods	Modality	Car (IOU=0.5)			Pedestrian (IOU=0.25)			Cyclist (IOU=0.25)		
			Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
0 (W/o Fog)	RTNH	4DR	35.8	41.7	35.9	42.0	38.6	34.4	75.3	69.8	63.2
	PointPillars	L	84.9	73.5	67.5	62.7	58.4	53.4	85.5	79.0	72.7
	InterFusion	L+4DR	67.6	65.8	58.8	73.7	70.1	64.7	90.3	87.0	81.2
	L4DR	L+4DR	85.0	76.6	69.4	**74.4**	**72.3**	**65.7**	**93.4**	**90.4**	83.0
	3D-LRF	L+4DR	85.4	74.9	69.2	69.0	67.9	62.7	92.0	87.2	81.7
	REL	L+4DR	**86.9**	**77.2**	**69.7**	74.2	70.4	65.1	92.8	90.1	**84.5**
2	PointPillars	L	67.0	51.4	44.4	53.1	47.2	42.7	69.6	62.7	57.2
	InterFusion	L+4DR	56.0	48.5	41.5	63.2	57.8	52.9	77.3	71.1	66.2
	L4DR	L+4DR	68.5	56.4	49.3	63.1	59.9	55.1	82.7	70.8	70.7
	3D-LRF	L+4DR	68.2	56.0	50.6	53.9	53.4	47.5	76.9	72.1	67.1
	REL	L+4DR	**68.5**	**64.4**	**56.9**	**65.5**	**62.1**	**57.1**	**86.6**	**81.6**	**76.1**
4	PointPillars	L	13.0	8.77	7.19	10.6	12.9	11.3	6.15	4.89	4.57
	InterFusion	L+4DR	15.2	10.8	8.40	25.7	25.1	22.6	6.68	7.95	6.99
	L4DR	L+4DR	26.9	26.2	21.6	33.1	30.7	27.9	30.3	29.7	26.3
	3D-LRF	L+4DR	25.0	20.6	16.4	25.2	28.1	24.2	21.1	20.2	18.4
	REL	L+4DR	**43.8**	**49.7**	**42.4**	**48.1**	**46.6**	**40.7**	**83.2**	**76.4**	**69.3**

Table 3: Quantitative results of different methods on Vod-Fog dataset under various fog levels. Best in **bold**, second in underline.

perior performance of REL on the K-Radar dataset V1.0. Specifically, our method surpasses 3D-LRF by 9.3% in total AP_{3D}, with consistent improvements under various adverse weather conditions. The superior performance of REL validates the rationality of generating virtual LiDAR points from weather-robust 4D radar data, demonstrating that 4D radar can effectively enhance the geometric structural information of LiDAR under adverse weather conditions. As shown in Figure 3, REL accurately predicts 3D bounding boxes under various weather conditions. In contrast, 3D-LRF, which relies heavily on LiDAR geometric structure, frequently produces false positives and missed detections in adverse weather scenarios.

On the K-Radar V2.0 dataset, we further validate the robustness of our method in complex scenarios and multi-category detection. As shown in Table 2, we compare methods based on LiDAR and 4D radar fusion, and all results under adverse weather conditions are aggregated for clarity. Compared with 3D-LRF, REL demonstrates significant performance gains in adverse weather, with AP_{3D} gains of 2% and 7.8% for the "Sedan" and "Bus or Truck" classes, respectively. These results highlight the effectiveness of virtual LiDAR points generated by 4D radar in mitigating the degradation of LiDAR point clouds in adverse conditions, particularly enhancing detection for larger objects where point degradation is more severe. Additionally, we integrate REL into the L4DR framework to test the generalization of our method. Compared with the original L4DR, consistent improvements are observed across object categories, with AP_{3D} gains of 1.3% for the "Sedan" and 3.9% gain for the "Bus or Truck" class, underscoring its adaptability across different LiDAR-4D radar fusion frameworks.

Results on Vod-Fog Dataset

On the Vod-Fog dataset, we compare the performance of REL with other LiDAR-4D radar methods under varying levels of fog density, as detailed in Table 3. Under clear conditions (fog level = 0), all methods exhibit good detection performance. However, as fog density increases, the quality of LiDAR point clouds degrades. Methods such as InterFusion (Wang et al. 2022b), 3D-LRF (Chae, Kim, and Yoon 2024), and L4DR (Huang et al. 2025b) show substantial performance drops due to their heavy reliance on geometric structures provided by the contaminated LiDAR data. In contrast, our method addresses this issue by dynamically generating virtual LiDAR points guided by 4D radar data, which effectively enhances the geometric structures of degraded point clouds, enabling REL to maintain relatively stable detection performance even in foggy conditions.

To further investigate the source of the performance gain, we include a 4D radar only baseline in Table 3. Under the most severe foggy conditions (fog level = 4), the 4D radar only method outperforms L4DR by 6.5–45.0 mAP but still underperforms our REL approach by 6.1–8.0 mAP. This indicates that previous methods struggle with negative fusion due to low-quality LiDAR in adverse weather, while our method effectively handles fusion in these conditions.

Ablation Study

Effect of Each Component. We assess the detection accuracy, inference speed (FPS), and computational cost (GFLOPs) of each REL module on K-Radar V1.0, as shown in Table 4. The first row represents the baseline model that uses the original LiDAR point clouds and 4D radar for 3D fusion. The second row introduces VLG to generate virtual LiDAR points. Both AP_{BEV} and AP_{3D} show improvements, validating the effectiveness of leveraging robust 4D radar data to supplement LiDAR geometric structure. The third row incorporates PGCA to enhance the features of virtual points, further boosting detection performance. The results show that PGCA effectively resolves the feature misalignment between virtual and real points. The fourth row integrates AFF, which effectively fuses the features of virtual LiDAR points and real LiDAR points through an adaptive weighting strategy, delivering the most optimal performance (i.e., our final model). Although REL introduces modest increases in runtime and computational cost than the baseline, it achieves significantly higher detection accuracy.

Module			IOU=0.5			
VLG	PGCA	AFF	AP_{BEV}	AP_{3D}	FPS	GFLOPs
			72.2	46.0	24.1	125.3
✓			74.1	50.4	23.7	127.7
✓	✓		75.1	52.5	22.4	130.2
✓	✓	✓	**75.6**	**54.5**	20.7	130.6

Table 4: Ablation of REL components on K-Radar V1.0.

Impact of Module Deployment at Different Stages. We evaluate the effects of deploying the proposed VLG-PGCA-AFF modules at different stages of the baseline on K-Radar V1.0, as shown in Table 5. The baseline has a three-level downsampling structure to extract multi-scale features from LiDAR points and 4D radar data, with an AP_{3D} of 46.0%. Notably, integrating our modules at any stage improves performance, though the magnitude varies across stages.

Since the generation of virtual points depends on 4D radar features, integrating these modules in Stage I suffers from a shallow network architecture that leads to insufficient noise suppression in the extracted 4D radar features. This results in relatively poor-quality virtual points. Stage II strikes a balance between LiDAR sparsity and radar feature representation, thus producing more accurate virtual points that enhance the geometric structure. However, in Stage III, repeated downsampling causes high sparsity in LiDAR points and a reduction in the number of 4D radar points, yielding inferior performance compared to Stage II.

Based on this analysis, we choose to deploy the proposed modules in both Stage-II and Stage-III. In this way, Stage-II generates virtual points to supplement the geometric structure of LiDAR and provides fused features as input to the deep LiDAR branch. Subsequently, Stage-III further refines the geometric structure of LiDAR. This dual-stage deployment effectively avoids interference from shallow noise while addressing the loss of deep structural information. Its optimal performance is validated through experiments testing all possible combinations.

Stage	IOU=0.5	
	$AP_{BEV}(\uparrow)$	$AP_{3D}(\uparrow)$
I	74.1	49.5
II	75.5	51.6
III	74.9	48.8
I + II + III	75.5	51.3
I + III	75.2	49.7
I + II	74.7	50.9
II + III	**75.6**	**54.5**

Table 5: Performance comparison of proposed modules deployment at different stages, tested on K-Radar V1.0.

Conclusion

In this paper, we address the problem of reduced 3D object detection accuracy caused by the degradation of LiDAR point geometric structure in adverse weather conditions. We propose REL, a 4D radar-guided LiDAR geometric enhancement method. By innovatively generating virtual LiDAR points from 4D radar data, REL effectively improves the spatial density and geometric structure of degraded LiDAR point clouds. Moreover, the proposed PGCA enhances the virtual LiDAR point features, while AFF performs weighted fusion of virtual and real LiDAR features in voxel space. Extensive experiments on K-Radar and Vod-Fog datasets demonstrate the significant improvements achieved by REL under various adverse weather conditions. Additionally, our method seamlessly integrates with existing LiDAR-4D radar fusion frameworks, offering an effective geometric enhancement solution for 3D object detection in challenging weather conditions.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grants 62222302, U2441251 and U22B2055, and in part by the Fundamental Research Funds for the Central Universities under Grant FRF-TP-22-003C1

References

Bang, G.; Choi, K.; Kim, J.; Kum, D.; and Choi, J. W. 2024. Radardistill: Boosting radar-based object detection performance via knowledge distillation from lidar features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15491–15500.

Chae, Y.; Kim, H.; and Yoon, K.-J. 2024. Towards robust 3d object detection with lidar and 4d radar fusion in various weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15162–15172.

Fan, L.; Wang, J.; Chang, Y.; Li, Y.; Wang, Y.; and Cao, D. 2024. 4D mmWave radar for autonomous driving perception: a comprehensive survey. *IEEE Transactions on Intelligent Vehicles*, 7(4): 4606–4620.

Gu, Z.; Ma, J.; Huang, Y.; Wei, H.; Chen, Z.; Zhang, H.; and Hong, W. 2025. HGSFusion: Radar-camera fusion with hybrid generation and synchronization for 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3, 3185–3193.

Gupta, A.; Jain, S.; Choudhary, P.; and Parida, M. 2024. Dynamic object detection using sparse LiDAR data for autonomous machine driving and road safety applications. *Expert Systems with Applications*, 255: 124636–124563.

Hahner, M.; Sakaridis, C.; Dai, D.; and Van Gool, L. 2021. 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*, 15283–15292.

Hoang, H. A.; Bui, D. C.; and Yoo, M. 2024. TSST-Det: transformation-based 3-D object detection via a spatial shape transformer. *IEEE Sensors Journal*, 24(5): 7126–7139.

Huang, H.; Yan, X.; Yang, J.; Cao, Y.; and Zhang, X. 2023. LIDSOR: A filter for removing rain and snow noise points from LiDAR point clouds in rainy and snowy weather. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48: 733–740.

Huang, X.; Wang, J.; Xia, Q.; Chen, S.; Yang, B.; Li, X.; Wang, C.; and Wen, C. 2025a. V2X-R: Cooperative LiDAR-4D Radar Fusion with Denoising Diffusion for 3D Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27390–27400.

Huang, X.; Wu, H.; Li, X.; Fan, X.; Wen, C.; and Wang, C. 2024. Sunshine to rainstorm: Cross-weather knowledge distillation for robust 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3, 2409–2416.

Huang, X.; Xu, Z.; Wu, H.; Wang, J.; Xia, Q.; Xia, Y.; Li, J.; Gao, K.; Wen, C.; and Wang, C. 2025b. L4dr: Lidar-4dradar fusion for weather-robust 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 4, 3806–3814.

Jiang, T.; Zhuang, L.; An, Q.; Wang, J.; Xiao, K.; and Wang, A. 2022. T-RODNet: Transformer for vehicular millimeter-wave radar object detection. *IEEE Transactions on Instrumentation and Measurement*, 72: 1–12.

Kong, H.; Huang, C.; Yu, J.; and Shen, X. 2024. A survey of mmwave radar-based sensing in autonomous vehicles, smart homes and industry. *IEEE Communications Surveys & Tutorials*, 27(1): 463–508.

Kong, S.-H.; Paek, D.-H.; and Cho, S. J. 2023. RTNH+: Enhanced 4D Radar Object Detection Network using Combined CFAR-based Two-level Preprocessing and Vertical Encoding. *CoRR*, 1–14.

Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12697–12705.

Li, P.; Wang, P.; Berntorp, K.; and Liu, H. 2022. Exploiting temporal relations on radar perception for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17071–17080.

Liu, C.; Kelber, F.; Vogginger, B.; and Mayr, C. 2024a. Cacfar is convolution: Fast target detection with machine learning accelerator. In *Mediterranean Conference on Embedded Computing*, 1–6.

Liu, J.; Zhao, Q.; Xiong, W.; Huang, T.; Han, Q.-L.; and Zhu, B. 2023. SMURF: Spatial multi-representation fusion for 3D object detection with 4D imaging radar. *IEEE Transactions on Intelligent Vehicles*, 9(1): 799–812.

Liu, Y.; Tian, Y.; Sun, B.; Wang, Y.; and Wang, F.-Y. 2022. Parallel lidars meet the foggy weather. *IEEE Journal of Radio Frequency Identification*, 6: 867–870.

Liu, Y.; Zhang, Y.; Lan, R.; Cheng, C.; and Wu, Z. 2025. AWARDistill: Adaptive and robust 3D object detection in adverse conditions through knowledge distillation. *Expert Systems with Applications*, 266: 126032–126049.

Liu, Z.; Hou, J.; Wang, X.; Ye, X.; Wang, J.; Zhao, H.; and Bai, X. 2024b. Lion: Linear group rnn for 3d object detection in point clouds. *Advances in Neural Information Processing Systems*, 37: 13601–13626.

Ma, T.; Yang, X.; Zhou, H.; Li, X.; Shi, B.; Liu, J.; Yang, Y.; Liu, Z.; He, L.; Qiao, Y.; et al. 2023. Detzero: Rethinking offboard 3d object detection with long-term sequential point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6736–6747.

Matteazzi, A.; Colling, P.; Arnold, M.; and Tutsch, D. 2025. Adverse Weather Conditions Augmentation of LiDAR Scenes with Latent Diffusion Models. *CoRR*, 1–5.

Meyer, M.; Kuschk, G.; and Tomforde, S. 2021. Graph convolutional networks for 3d object detection on radar data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3060–3069.

Naich, A. Y.; and Carrión, J. R. 2024. LiDAR-based intensity-aware outdoor 3D object detection. *Sensors*, 24(9): 2942–2959.

Oh, M.; and Dong-O'Brien, J. 2025. Assessing the safety impacts of winter road maintenance operations using connected vehicle data. *Accident Analysis & Prevention*, 209–220: 107837.

Paek, D.-H.; Kong, S.-H.; and Wijaya, K. T. 2022. K-radar: 4d radar object detection for autonomous driving in various weather conditions. *Advances in Neural Information Processing Systems*, 35: 3819–3829.

Paek, D.-H.; Kong, S.-H.; and Wijaya, K. T. 2023. Enhanced k-radar: Optimal density reduction to improve detection performance and accessibility of 4d radar tensor-based object detection. In *IEEE Intelligent Vehicles Symposium*, 1–6.

Palffy, A.; Pool, E.; Baratam, S.; Kooij, J. F.; and Gavrila, D. M. 2022. Multi-class road user detection with 3+ 1d radar in the view-of-delft dataset. *IEEE Robotics and Automation Letters*, 7(2): 4961–4968.

Qi, Y.; Liu, C.; Scaioni, M.; Li, Y.; Qiao, Y.; Ma, X.; Wu, H.; Zhang, K.; and Wang, D. 2024. Geometric information constraint 3D object detection from LiDAR point cloud for autonomous vehicles under adverse weather. *Transportation research part C: emerging technologies*, 161–183: 104555.

Roriz, R.; Campos, A.; Pinto, S.; and Gomes, T. 2021. DIOR: A hardware-assisted weather denoising solution for LiDAR point clouds. *IEEE Sensors Journal*, 22(2): 1621–1628.

Scheuble, D.; Linnhoff, C.; Bijelic, M.; Elster, L.; Rosenberger, P.; Ritter, W.; and Winner, H. 2024. Simulating Road Spray Effects in Automotive Lidar Sensor Models. In *IEEE Intelligent Vehicles Symposium*, 659–666.

Song, J.; Zhao, L.; and Skinner, K. A. 2024. Lirafusion: Deep adaptive lidar-radar fusion for 3d object detection. In *IEEE International Conference on Robotics and Automation*, 18250–18257.

Song, Z.; Liu, L.; Jia, F.; Luo, Y.; Jia, C.; Zhang, G.; Yang, L.; and Wang, L. 2024. Robustness-aware 3d object detection in autonomous driving: A review and outlook. *IEEE Transactions on Intelligent Transportation Systems*, 25(11): 15407–15436.

Teufel, S.; Volk, G.; Von Bernuth, A.; and Bringmann, O. 2022. Simulating realistic rain, snow, and fog variations for comprehensive performance characterization of LiDAR perception. In *IEEE 95th Vehicular Technology Conference:(VTC2022-Spring)*, 1–7.

Wang, L.; Zhang, X.; Li, J.; Xv, B.; Fu, R.; Chen, H.; Yang, L.; Jin, D.; and Zhao, L. 2022a. Multi-modal and multi-scale fusion 3D object detection of 4D radar and LiDAR for autonomous driving. *IEEE Transactions on Vehicular Technology*, 72(5): 5628–5641.

Wang, L.; Zhang, X.; Xv, B.; Zhang, J.; Fu, R.; Wang, X.; Zhu, L.; Ren, H.; Lu, P.; Li, J.; et al. 2022b. InterFusion: Interaction-based 4D radar and LiDAR fusion for 3D object detection. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 12247–12253.

Wang, W.; You, X.; Chen, L.; Tian, J.; Tang, F.; and Zhang, L. 2022c. A scalable and accurate de-snowing algorithm for LiDAR point clouds in winter. *Remote Sensing*, 14(6): 1468–1485.

Wozniak, M. K.; Hansson, M.; Thiel, M.; and Jensfelt, P. 2024. Uada3d: Unsupervised adversarial domain adaptation for 3d object detection with sparse lidar and large domain gaps. *IEEE Robotics and Automation Letters*, 9(12): 11210–11217.

Xu, B.; Zhang, X.; Wang, L.; Hu, X.; Li, Z.; Pan, S.; Li, J.; and Deng, Y. 2021. RPFA-Net: A 4D radar pillar feature attention network for 3D object detection. In *IEEE International Intelligent Transportation Systems Conference*, 3061–3066.

Yataka, R.; Wang, P.; Boufounos, P.; and Takahashi, R. 2024. SIRA: scalable inter-frame relation and association for radar perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15024–15034.

Zhang, A.; Nowruzi, F. E.; and Laganiere, R. 2021. Raddet: Range-azimuth-doppler based radar object detection for dynamic road users. In *18th Conference on Robots and Vision*, 95–102.

Zhang, C.; Wang, H.; Cai, Y.; Chen, L.; and Li, Y. 2024a. TransFusion: multi-modal robust fusion for 3D object detection in foggy weather based on spatial vision transformer. *IEEE Transactions on Intelligent Transportation Systems*, 25(9): 10652–10666.

Zhang, G.; Fan, L.; He, C.; Lei, Z.; ZHANG, Z.-X.; and Zhang, L. 2024b. Voxel mamba: Group-free state space models for point cloud based 3d object detection. *Advances in Neural Information Processing Systems*, 37: 81489–81509.

Zhang, M.; Abdulatif, S.; Loesch, B.; Altmann, M.; Schwarz, M.; and Yang, B. 2024c. Exploring domain shift on radar-based 3D object detection amidst diverse environmental conditions. In *IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*, 3513–3518.

Zhang, Y.; Ding, M.; Yang, H.; Niu, Y.; Feng, Y.; Ohtani, K.; and Takeda, K. 2023. L-DIG: A GAN-based method for LiDAR point cloud processing under snow driving conditions. *Sensors*, 23(21): 8660–8667.

Zhang, Y.; Ding, M.; Yang, H.; Niu, Y.; Ge, M.; Ohtani, K.; Zhang, C.; and Takeda, K. 2024d. LiDAR Point Cloud Augmentation for Adverse Conditions Using Conditional Generative Model. *Remote Sensing*, 16(12): 2247–2268.

Zhao, H.; Jiang, L.; Jia, J.; Torr, P. H.; and Koltun, V. 2021. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16259–16268.

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.