

# L4DR: LiDAR-4DRadar Fusion for Weather-Robust 3D Object Detection

Xun Huang<sup>1,2,3</sup>, Ziyu Xu<sup>1,2</sup>, Hai Wu<sup>1,2</sup>, Jinlong Wang<sup>1,2</sup>, Qiming Xia<sup>1,2</sup>,  
Yan Xia<sup>4</sup>, Jonathan Li<sup>5</sup>, Kyle Gao<sup>5</sup>, Chenglu Wen<sup>1,2\*</sup>, Cheng Wang<sup>1,2</sup>

<sup>1</sup>Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen University, China

<sup>2</sup>Key Laboratory of Multimedia Trusted Perception and Efficient Computing,  
Ministry of Education of China, Xiamen University, China

<sup>3</sup>Zhongguancun Academy, China

<sup>4</sup>Technische Universität München, Germany

<sup>5</sup>University of Waterloo, Canada

## Abstract

LiDAR-based 3D object detection is crucial for autonomous driving. However, due to the quality deterioration of LiDAR point clouds, it suffers from performance degradation in adverse weather conditions. Fusing LiDAR with the weather-robust 4D radar sensor is expected to solve this problem; however, it faces challenges of significant differences in terms of data quality and the degree of degradation in adverse weather. To address these issues, we introduce L4DR, a weather-robust 3D object detection method that effectively achieves LiDAR and 4D Radar fusion. Our L4DR proposes **Multi-Modal Encoding (MME)** and **Foreground-Aware Denoising (FAD)** modules to reconcile sensor gaps, which is the first exploration of the complementarity of early fusion between LiDAR and 4D radar. Additionally, we design an **Inter-Modal and Intra-Modal ( $\{IM\}^2$ )** parallel feature extraction backbone coupled with a **Multi-Scale Gated Fusion (MSGF)** module to counteract the varying degrees of sensor degradation under adverse weather conditions. Experimental evaluation on a VoD dataset with simulated fog proves that L4DR is more adaptable to changing weather conditions. It delivers a significant performance increase under different fog levels, improving the 3D mAP by up to 20.0% over the traditional LiDAR-only approach. Moreover, the results on the K-Radar dataset validate the consistent performance improvement of L4DR in real-world adverse weather conditions.

**Code** — <https://github.com/ylwhxht/L4DR>

**Extended version** — <https://arxiv.org/abs/2408.03677>

## Introduction

3D object detection is a fundamental vision task of unmanned platforms, extensively utilized in applications such as intelligent robot navigation (Ghita et al. 2024; Xu et al. 2022) and autonomous driving (Bijelic et al. 2020). For example, Full Driving Automation (FDA, Level 5) relies on weather-robust 3D object detection, which provides precise 3D bounding boxes even under various challenging adverse weather conditions (Qian et al. 2021). Owing to the high resolution and

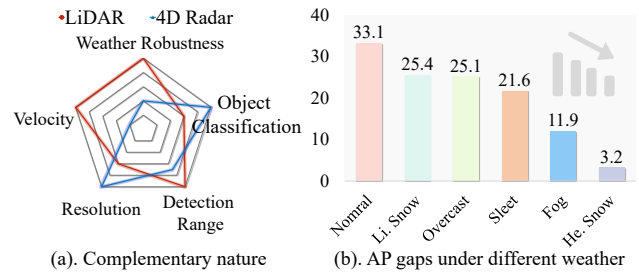


Figure 1: (a) The radar chart illustrates the complementary nature of LiDAR and 4D radar sensors. (b) AP gaps between LiDAR and 4D radar in real different weather (the extent to which LiDAR is superior to 4D radar).

strong interference resistance of LiDAR sensors, LiDAR-based 3D object detection has emerged as a mainstream area of research (Wu et al. 2024; Xia et al. 2024; Deng et al. 2025). However, LiDAR sensors exhibit considerable sensitivity to weather conditions. In adverse scenarios, the scanned LiDAR point clouds suffer from substantial degradation and increased noise (Huang et al. 2024a; Hahner et al. 2021). This degradation negatively impacts 3D detectors, compromising the reliability of autonomous perception systems.

Aside from LiDAR, 4D (range, azimuth, Doppler, and elevation) millimeter-wave radar is gaining recognition as a vital perception sensor (Han et al. 2023; Sun and Zhang 2021; Li, Qian, and Xia 2024). As shown in Fig.1 (a), 4D radar outperforms LiDAR in weather robustness, velocity measurement, and detection range. The millimeter-wave signals of 4D radar have wavelengths much larger than the tiny particles in fog, rain, and snow (Golovachev et al. 2018; Balal, Pinhasi, and Pinhasi 2016), exhibiting reduced susceptibility to weather disturbances. As shown in Fig.1 (b), the performance gap between LiDAR and 4D radar decreases as the severity of the weather rises. Hence, the 4D radar sensor is more suitable for adverse weather conditions than the LiDAR sensor. However, LiDAR has advantages in terms of higher distance and angular point resolution. These circumstances make it feasible to promote the full fusion of 4D radar and LiDAR data to improve 3D object detection. Pioneering approaches such as InterFusion (Wang et al. 2022b), M<sup>2</sup>Fusion (Wang

\*Corresponding author. Email: clwen@xmu.edu.cn.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

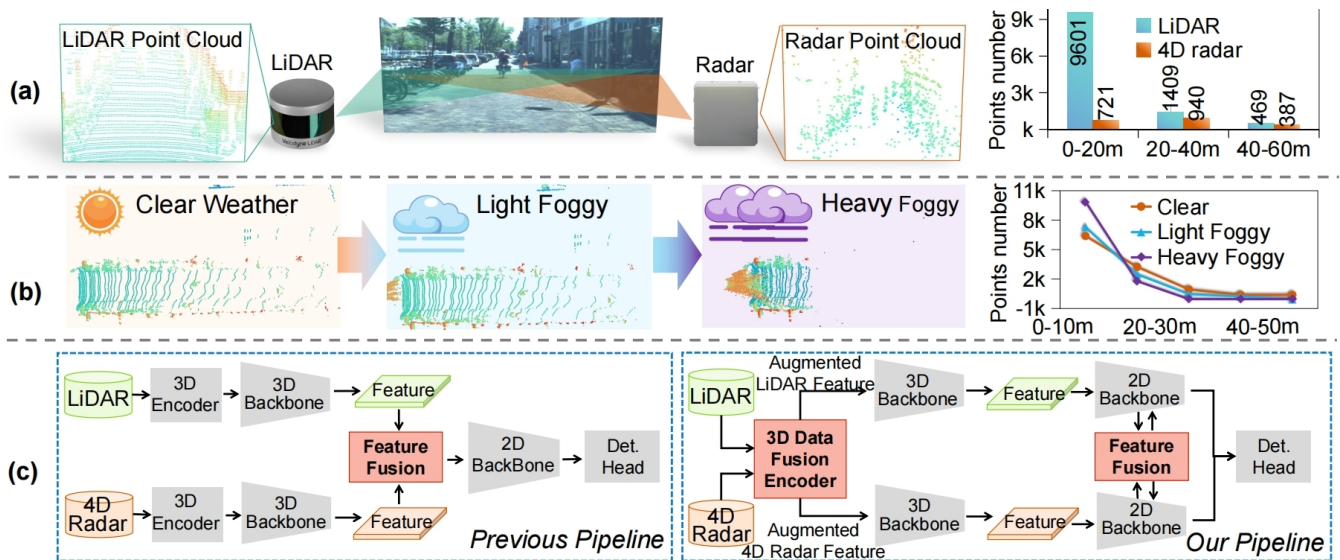


Figure 2: (a) Significant quality disparity between the LiDAR and the 4D radar. (b) Severe degradation of LiDAR data quality in adverse weather. (c) Comparison of previous LiDAR-4DRadar fusion and our fusion, highlighting our innovative framework designs to address challenges (a) and (b).




Figure 3: Performance comparison of our L4DR and LiDAR-only in (a) various simulated fog levels (FL denotes fog level) and (b) real-world adverse weather.

et al. 2022a), and 3D-LRF (Chae, Kim, and Yoon 2024) conducted initial attempts to fuse LiDAR and 4D radar data, have demonstrated significant performance improvements over single-sensor models.

Despite progress, the significant data quality disparities and varying sensor degradation in adverse weather conditions remain largely unaddressed. As shown in Fig.2 (a), a primary challenge arises from the significant quality disparity between the LiDAR and 4D radar sensors. A second challenge pertains to varying sensor degradation under adverse weather conditions. Fig.2 (b) shows LiDAR sensors undergo severe data degradation in adverse weather. Conversely, the data quality decrease in 4D radar is significantly lower (Han et al. 2023; Sun and Zhang 2021) (Details are given in extended version (Huang et al. 2024b)). The varying degradation of sensors leads to fluctuations in the fused features, resulting in difficulty for the model to maintain high performance under adverse weather conditions. Therefore, developing a weather-

robust 3D object detector requires overcoming the challenges of significant quality disparities and varying degradation.

To address the above challenges, we propose L4DR, a novel two-stage fusion framework for LiDAR and 4D radar (shown in Fig. 2 (c)). The first-stage fusion is performed by the 3D Data Fusion Encoder, which contains the Foreground-Aware Denoising (FAD) module and the Multi-Modal Encoder (MME) module. This first-stage data fusion augments LiDAR and 4D radar with each other to tackle the challenge of significant data quality disparities between LiDAR and 4D radar. The second-stage Feature Fusion is to fuse features between 2D Backbones to address the second challenge of the varying sensor degradation. It consists of the Inter-Modal and Intra-Modal ( $\{IM\}^2$ ) backbone and the Multi-Scale Gated Fusion (MSGF) module. Different from traditional methods that fuse features before extraction on the left side of Fig.2 (c), L4DR integrates a continuous fusion process throughout the feature extraction, adaptively focusing on the significant modal features in different weather conditions.

In Fig. 3, our comprehensive experimental results showcase the superior performance of L4DR across various simulated and real-world adverse weather disturbances. Our main contributions are as follows:

- We introduce the innovative Multi-Modal Encoder (MME) module, which achieves LiDAR and 4D radar data early fusion without resorting to error-prone processes (e.g., depth estimation). It effectively bridges the substantial LiDAR and 4D radar data quality disparities.
- We designed an  $\{IM\}^2$  backbone with a Multi-Scale Gated Fusion (MSGF) module, adaptively extracting salient features from LiDAR and 4D radar in different weather conditions. This enables the model to adapt to varying levels of sensor degradation under adverse weather conditions.

- Extensive experiments on the two benchmarks, VoD and K-Radar, demonstrate the effectiveness of our L4DR under various levels and types of adverse weather, achieving new state-of-the-art performances on both datasets.

## Related Work

**LiDAR-based 3D object detection.** Researchers have developed single-stage and two-stage methods to tackle challenges for 3D object detection. Single-stage detectors such as VoxelNet (Yan, Mao, and Li 2018), PointPillars (Lang et al. 2019), 3DSSD (Yang et al. 2020), DSVT (Wang et al. 2023a) utilize PointNet++ (Qi et al. 2017), sparse convolution, or other point feature encoder to extract features from point clouds and perform detection in the Bird’s Eye View (BEV) space. Conversely, methods such as PV-RCNN (Shi et al. 2020), PV-RCNN++ (Shi et al. 2022), Voxel-RCNN (Deng et al. 2021), and VirConv (Wu et al. 2023) focus on two-stage object detection, integrating RCNN networks into 3D detectors. Even though these mainstream methods have gained excellent performance in normal weather, they still lack robustness under various adverse weather conditions.

**LiDAR-based 3D object detection in adverse weather.** LiDAR sensors may undergo degradation under adverse weather conditions. Physics-based weather simulations (Hahner et al. 2022, 2021; Kilic et al. 2021) have been explored to produce point clouds under adverse weather to alleviate data scarcity. (Charron, Phillips, and Waslander 2018; Heinzler et al. 2020) utilized the DROR, DSOR, or CNN to classify and filter LiDAR noise points. (Xu et al. 2021) designed a general completion framework that addresses the domain gaps across different weather conditions. (Huang et al. 2024a) designed a general knowledge distillation framework that transfers sunny performance to rainy performance. However, these methods primarily rely on single-LiDAR modal data, which will be constrained by the decline in the quality of LiDAR under adverse weather conditions.

**LiDAR-radar fusion-based 3D object detection.** MVDNet (Qian et al. 2021) has designed a framework for fusing LiDAR and radar. ST-MVDNet (Li et al. 2022) and ST-MVDNet++ (Li, O’Toole, and Kitani 2023) incorporate a self-training teacher-student to MVDNet to enhance the model. Bi-LRFusion (Wang et al. 2023b) employs a bidirectional fusion strategy to improve dynamic object detection performance. However, these studies only focus on 3D radar and LiDAR. As research progresses, the newest studies continue to drive the development of LiDAR-4D radar fusion. M<sup>2</sup>-Fusion (Wang et al. 2022a), InterFusion (Wang et al. 2022b), and 3D-LRF (Chae, Kim, and Yoon 2024) explore LiDAR and 4D radar fusion. However, these methods have not considered and overcome the challenges of fusing 4D radar and LiDAR under adverse weather conditions.

## Methodology

### Problem Statement and Overall Design

**LiDAR-4D radar fusion-based 3D object detection.** For an outdoor scene, we denote LiDAR point cloud as  $\mathcal{P}^l =$

$\{p_i^l\}_{i=1}^{N_l}$  and 4D radar point cloud as  $\mathcal{P}^r = \{p_i^r\}_{i=1}^{N_r}$ , where  $p$  denotes 3D points. Subsequently, a multi-modal model  $\mathcal{M}$  extract deep features  $\mathcal{F}^m$  from  $\mathcal{P}^m$ , written as  $\mathcal{F}^m = g(f_{\mathcal{M}}(\mathcal{P}^m; \Theta))$ , where  $m \in \{l, r\}$ . The fusion features are then obtained by  $\mathcal{F}^f = \phi(\mathcal{F}^l, \mathcal{F}^r)$ , where  $\phi$  donates fusion method. The objective of 3D object detection is to regress the 3D bounding boxes  $B = \{b_i\}_{i=1}^{N_b}$ ,  $B \in \mathbb{R}^{N_b \times 7}$ .

**Significant data quality disparity.** As mentioned before, there is a dramatic difference between  $\mathcal{P}^l$  and  $\mathcal{P}^r$  in the same scene. To fully fuse two modalities, we can use  $\mathcal{P}^l$  to enhance the highly sparse  $\mathcal{P}^r$  that lacks discriminative details. Therefore, our L4DR includes a Multi-Modal Encoder (MME, Figure 4 (b)), which performs data early-fusion complementarity at the encoder. However, we found that direct data fusion would also cause noises in  $\mathcal{P}^r$  to spread to  $\mathcal{P}^l$ . Therefore, we integrated Foreground-Aware Denoising (FAD, Figure 4 (a)) into L4DR before MME to filter out most of the noise in  $\mathcal{P}^r$ .

**Varying sensor degradation in adverse weather.** Compared to 4D radar, the quality of LiDAR point cloud  $\mathcal{P}^l$  is more easily affected by adverse weather conditions, leading to varying feature presentations  $\mathcal{F}^l$ . Previous backbones focusing solely on fused inter-modal features  $\mathcal{F}^f$  overlook the weather robustness of 4D radar, leading to challenges in addressing the frequent fluctuation of  $\mathcal{F}^l$ . To ensure robust fusion across diverse weather conditions, we introduce the Inter-Modal and Intra-Modal ( $\{\text{IM}\}^2$ , Figure 4 (c)) backbone. This design simultaneously focuses on inter-modal and intra-modal features, enhancing model adaptability. However, redundancy between these features arises. Inspired by gated fusion techniques (Hosseinpour, Samadzadegan, and Javan 2022; Song, Zhao, and Skinner 2024a), we propose the Multi-Scale Gated Fusion (MSGF, Figure 4 (d)) module. MSGF utilizes inter-modal features  $\mathcal{F}^f$  to filter intra-modal features  $\mathcal{F}^l$  and  $\mathcal{F}^r$ , effectively reducing feature redundancy.

### Foreground-Aware Denoising (FAD)

Due to multipath effects, 4D radar contains significant noise points. Despite applying the Constant False Alarm Rate (CFAR) algorithm to filter out noise during the data acquisition process, the noise level remains substantial. It is imperative to further reduce clutter noise in 4D radar data before early data fusion to avoid spreading noise. Considering the minimal contribution of background points to object detection, this work introduces point-level foreground semantic segmentation to 4D radar denoising, performing a Foreground-Aware Denoising. Specifically, we first utilize PointNet++ (Qi et al. 2017) combined with a segmentation head as  $\chi$ , to predict the foreground semantic probability  $\mathcal{S} = \chi(\mathcal{P}^r)$  for each point in the 4D radar. Subsequently, points with a foreground probability below a predefined threshold  $\tau$  are filtered out, that is  $\mathcal{P}_{new}^r = \{p_i^r | \mathcal{S}_i \geq \tau\}$ . FAD effectively filters out as many noise points as possible while preventing the loss of foreground points.

### Multi-Modal Encoder (MME)

Even following denoising using FAD, there remains a significant quality disparity between LiDAR and 4D radar due

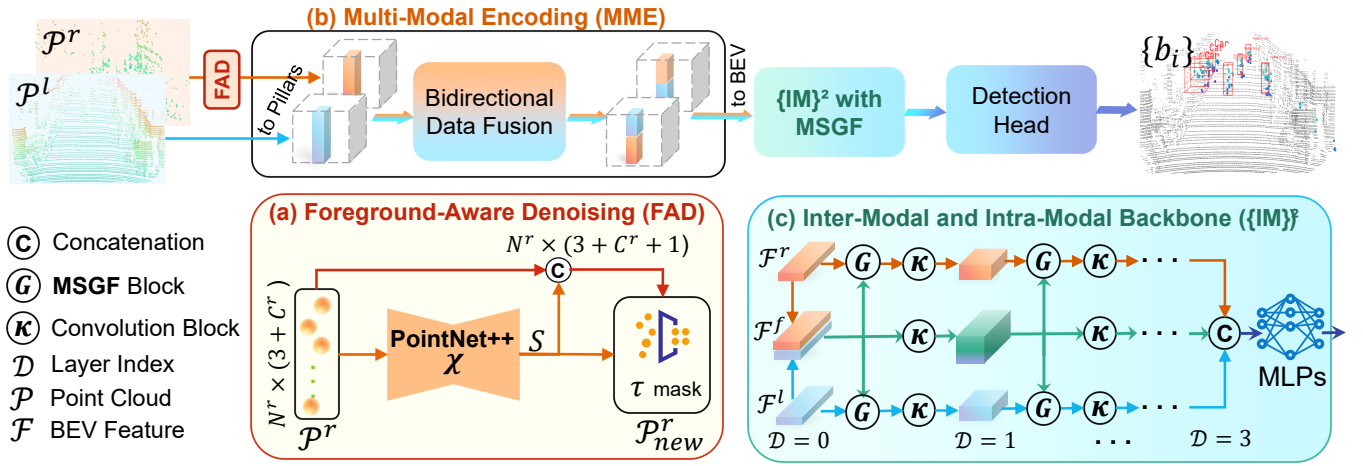


Figure 4: L4DR framework. (a) **Foreground-Aware Denoising (FAD)** performs denoising by segmenting foreground semantics per 4D radar point. Next, (b) **Multi-Modal Encoder (MME)** fuses bi-directional data for both LiDAR and 4D radar modalities at the Encoder stage to obtain higher quality BEV features. Finally, (c) **Inter-Modal and Intra-Modal ( $\{IM\}^2$ )** backbone coupled with **Multi-Scale Gated Fusion (MSGF)** uses a gating strategy to filter features to avoid redundant information while extracting inter-modal and intra-modal features in parallel.

to limitations in resolution. We thus design a Multi-Modal Encoder module that fuses LiDAR and radar points at an early stage to extract richer features.

As illustrated in Figure 5, we innovate the traditional unimodal Pillars coding into multimodal Pillars coding to perform initial fusion at the data level, extracting richer information for subsequent feature processing. Firstly, referring (Lang et al. 2019), we encode the LiDAR point cloud into a pillar set  $\mathcal{H}^l = \{h_i^l\}_{i=1}^N$ . Each LiDAR point  $p_{(i,j)}^l$  in pillar  $h_i^l$  encoded using encoding feature  $f_{(i,j)}^l$  as

$$f_{(i,j)}^l = [\mathcal{X}^l, \mathcal{Y}_{cl}^l, \mathcal{Z}^l, \lambda], \quad (1)$$

where  $\mathcal{X}^l = [x^l, y^l, z^l]$  is the coordinate of the LiDAR point,  $\mathcal{Y}_{cl}^l$  denotes the distance from the LiDAR point to the arithmetic mean of all LiDAR points in the pillar,  $\mathcal{Z}^l$  denotes the (horizontal) offset from the pillar center in  $x, y$  coordinates, and  $\lambda$  is the reflectance. Similarly, each 4D radar point  $p_{(i,k)}^r$  in pillar  $h_i^r$  encoded using encoding feature  $f_{(i,k)}^r$  as

$$f_{(i,k)}^r = [\mathcal{X}^r, \mathcal{Y}_{cr}^r, \mathcal{Z}^r, \mathcal{V}, \Omega], \quad (2)$$

where  $\mathcal{X}^r$ ,  $\mathcal{Y}_{cr}^r$ , and  $\mathcal{Z}^r$  are similar in meaning to those in Eq.1.  $\mathcal{V} = [\mathcal{V}_r, \mathcal{V}_a]$  is the relative and absolute radial velocity of Doppler information, and  $\Omega$  is Radar Cross-Section (RCS).

We then perform cross-modal feature propagation for LiDAR and radar pillar encoding features that occupy the same pillar coordinates. The fused LiDAR and 4D radar pillar encoding features  $\hat{f}_{(i,j)}^l$  and  $\hat{f}_{(i,k)}^r$  are obtained by fusing  $f_{(i,j)}^l$  and  $f_{(i,k)}^r$  as follows:

$$\begin{aligned} \hat{f}_{(i,j)}^l &= [\mathcal{X}^l, \mathcal{Y}_{cl}^l, \mathcal{Y}_{cr}^l, \mathcal{Z}^l, \lambda, \bar{\mathcal{V}}, \bar{\Omega}], \\ \text{and } \hat{f}_{(i,k)}^r &= [\mathcal{X}^r, \mathcal{Y}_{cl}^r, \mathcal{Y}_{cr}^r, \mathcal{Z}^r, \bar{\lambda}, \mathcal{V}, \Omega], \end{aligned} \quad (3)$$

where the overline denotes the average of all point features of another modality in that pillar.

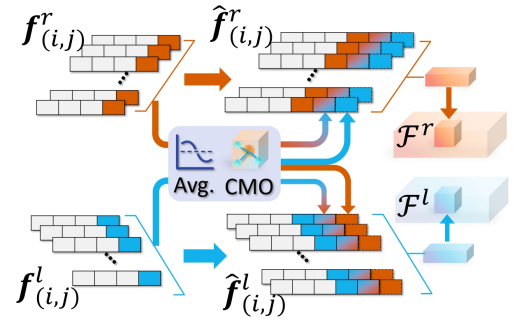


Figure 5: Bidirectional Data Fusion in MME. LiDAR-specific point features (blue) and radar-specific point features (red) located in the same pillar are averaged (Avg.) for feature propagation. And the cross-modal offsets (CMO) are computed to enrich the geometric features.

The feature propagation is beneficial because  $\lambda$  and  $\Omega$  are helpful for object classification, while Doppler information  $\mathcal{V}$  is crucial for distinguishing dynamic objects (Song, Zhao, and Skinner 2024b). Cross-modal feature sharing makes comprehensive use of these advantages and cross-modal offsets  $[\mathcal{Y}_{cl}^m, \mathcal{Y}_{cr}^m]$ ,  $m \in \{l, r\}$  also further enrich the geometric information. This MME method compensates for the data quality of 4D radar under normal weather conditions and can also enhance the quality of LiDAR in adverse weather. Subsequently, we applied a linear layer and max pooling operations to the fused pillar encoding features  $\hat{f}$  to obtain the corresponding modal BEV features  $\mathcal{F}$ .

### $\{IM\}^2$ Backbone and MSGF Block

To take full advantage of the respective advantages of LiDAR and 4D radar, it is necessary to focus on both inter-modal fea-

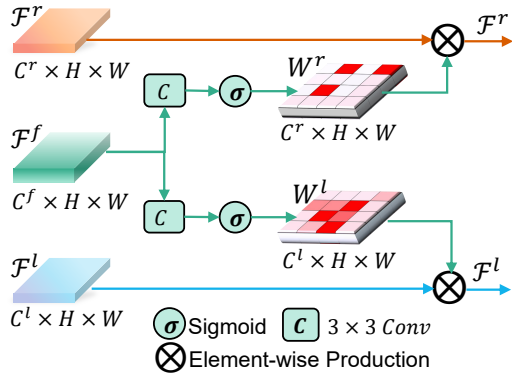


Figure 6: Gated fusion block in our MSGF. The gated block processes input features from LiDAR and 4D radar by using the fused inter-modal features. It generates adaptive gating weights from fused features and applies these weights via element-wise multiplication to filter out redundant information in intra-modal features.

tures and intra-modal features. We introduce the **Inter-Modal** and **Intra-Modal** backbone ( $\{\text{IM}\}^2$ ).  $\{\text{IM}\}^2$  serves as a multi-modal branch feature extraction module that concurrently extracts inter-modal feature ( $\mathcal{F}^f$ ) and intra-modal features ( $\mathcal{F}^l, \mathcal{F}^r$ ). Specifically, we fuse two Intra-Modal features to form an Inter-Modal feature,

$$\mathcal{F}^f = \phi(\mathcal{F}^l, \mathcal{F}^r), \quad (4)$$

where  $\phi$  denotes fusion approach (we use concatenation). Subsequently, we apply a convolutional block to each modal branch  $\mathcal{F}^l, \mathcal{F}^r$ , and  $\mathcal{F}^f$  independently,

$$\mathcal{F}_{\mathcal{D}}^m = \kappa(\mathcal{F}_{\mathcal{D}-1}^m), \quad (5)$$

where  $\mathcal{D} \in [1, 3]$  denotes layer, and  $m \in \{l, r, f\}$  indicates different modality.  $\kappa$  represents a convolutional layer with batch normalization and ReLU activation.

However, while  $\{\text{IM}\}^2$  addresses some deficiencies in feature representation, this naive approach inevitably introduces redundant features. Inspired by (Song, Zhao, and Skinner 2024b) to adaptively filter each modal feature, we design an MSGF for adaptive gated fusion on each LiDAR and 4D radar scale feature map.

As depicted in Fig. 6, the gated network  $\mathcal{G}$  in MSGF processes input feature maps from LiDAR  $\mathcal{F}^l$ , 4D radar  $\mathcal{F}^r$  and fused counterpart  $\mathcal{F}^f$ . Subsequently, on the LiDAR and 4D radar branches, the adaptive gating weights  $\mathcal{W}^l$  for  $\mathcal{F}^l$  and  $\mathcal{W}^r$  for  $\mathcal{F}^r$  were obtained by a convolution block and sigmoid activation function, respectively. These weights are applied to the initial feature via element-wise multiplication, thus enabling filter  $\mathcal{F}^l$  and  $\mathcal{F}^r$  in the gated mechanism. Formally, the gated network  $\mathcal{G}$  guides  $\mathcal{F}^l$  and  $\mathcal{F}^r$  at convolution layer index  $\mathcal{D}$  to filter out redundant information as following:

$$\mathcal{F}_{\mathcal{D}}^m = \mathcal{G}_{\mathcal{D}}(\mathcal{F}_{\mathcal{D}}^m, \mathcal{F}_{\mathcal{D}}^f), m \in \{l, r\}, \quad (6)$$

and  $\mathcal{G}_{\mathcal{D}}^m(\mathcal{F}_{\mathcal{D}}^m, \mathcal{F}_{\mathcal{D}}^f) = \mathcal{F}_{\mathcal{D}}^m * \delta(\kappa(\mathcal{F}_{\mathcal{D}}^f))$ ,

where  $\kappa$  is a  $3 \times 3$  convolution block and  $\delta$  a sigmoid function.  $\mathcal{F}^f$  is the fused feature with information about the interactions between modalities. It discerns whether features in  $\mathcal{F}^l$

and  $\mathcal{F}^r$  are helpful or redundant. Using  $\mathcal{F}^f$  for gated filtering can flexibly weight and extract features from  $\mathcal{F}^l$  and  $\mathcal{F}^r$  while significantly reducing feature redundancy.

## Loss Function and Training Strategy

We trained our L4DR with the following losses :

$$\mathcal{L}_{all} = \beta_{cls} \mathcal{L}_{cls} + \beta_{loc} \mathcal{L}_{loc} + \beta_{fad} \mathcal{L}_{fad}, \quad (7)$$

where  $\beta_{cls} = 1, \beta_{loc} = 2, \beta_{fad} = 0.5$ , the  $\mathcal{L}_{cls}$  is object classification focal loss, the  $\mathcal{L}_{loc}$  is object localization regression loss, and the  $\mathcal{L}_{fad}$  is focal classification loss of 4D radar foreground point in FAD module. We use Adam optimizer with lr =  $1e-3, \beta_1 = 0.9, \beta_2 = 0.999$ .

## Experiments

### Implement Details

We implement L4DR with PointPillars (Lang et al. 2019), the most commonly used base architecture in radar-based, LiDAR and 4D radar fusion-based 3D object detection. This can effectively verify the effectiveness of our L4DR and avoid unfair comparisons caused by inherent improvements in the base architecture. We set  $\tau$  in section 3.2 as 0.3 while training and 0.2 while inferring. We conduct all experiments with a batch size of 16 on 2 RTX 3090 GPUs. Other parameter settings refer to the default official configuration in the OpenPCDet (Team et al. 2020) tool.

### Dataset and Evaluation Metrics

**K-Radar dataset.** The K-Radar dataset (Paek, Kong, and Wijaya 2022) contains 58 sequences with 34944 frames of 64-line LiDAR, camera, and 4D radar data in various weather conditions. According to the official K-Radar split, we used 17458 frames for training and 17536 frames for testing. We adopt two evaluation metrics for 3D object detection:  $AP_{3D}$  and  $AP_{BEV}$  of the class "Sedan" at IoU = 0.5.

**View-of-Delft (VoD) dataset.** The VoD dataset (Palffy et al. 2022) contains 8693 frames of 64-line LiDAR, camera, and 4D radar data. Following the official partition, we divided the dataset into a training and validation set with 5139 and 1296 frames. All of the methods used the official radar with 5 scans accumulation and single frame LiDAR. Meanwhile, to explore the performance under different fog intensities, following a series of previous work (Qian et al. 2021; Li et al. 2022), we similarly performed fog simulations (Hahner et al. 2021) (with fog level  $\mathcal{L}$  from 0 to 4, fog density ( $\alpha$ ) = [0.00, 0.03, 0.06, 0.10, 0.20]) on the VoD dataset and kept 4D Radar unchanged to simulate various weather conditions. We named it the **Vod-Fog** dataset in the following. *Note-worthy, we used two evaluation metrics groups on the VoD dataset. The VoD official metrics are used to compare with previous state-of-the-art methods. The KITTI official metrics are used to analyze the performance of objects with different difficulties.*

### Results on K-Radar Adverse Weather Dataset

Following 3D-LRF (Chae, Kim, and Yoon 2024), we compare our L4DR with different modality-based 3D object detection methods: PointPillars (Lang et al. 2019), RTNH (Paek,

Methods	Modality	Metric	Total	Normal	Overcast	Fog	Rain	Sleet	Lightsnow	Heavysnow
RTNH (NeurIPS 2022)	4DR	$AP_{BEV}$	41.1	41.0	44.6	45.4	32.9	50.6	81.5	56.3
		$AP_{3D}$	37.4	37.6	42.0	41.2	29.2	49.1	63.9	43.1
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	<u>59.8</u>	28.2	<u>31.4</u>	50.7	24.6
InterFusion (IROS 2023)	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	9.33	<u>56.8</u>	25.7
3D-LRF (CVPR 2024)	L+4DR	$AP_{BEV}$	<u>73.6</u>	<u>72.3</u>	88.4	<u>86.6</u>	<u>76.6</u>	47.5	<u>79.6</u>	<b>64.1</b>
		$AP_{3D}$	<u>45.2</u>	<u>45.3</u>	<u>55.8</u>	51.8	<u>38.3</u>	23.4	<b>60.2</b>	<u>36.9</u>
L4DR (Ours)	L+4DR	$AP_{BEV}$	<b>77.5</b>	<b>76.8</b>	<b>88.6</b>	<b>89.7</b>	<b>78.2</b>	<b>59.3</b>	<b>80.9</b>	<u>53.8</u>
		$AP_{3D}$	<b>53.5</b>	<b>53.0</b>	<b>64.1</b>	<b>73.2</b>	<b>53.8</b>	<b>46.2</b>	52.4	<b>37.0</b>

Table 1: Quantitative results of different 3D object detection methods on K-Radar dataset. We present the modality of each method (L: LiDAR, 4DR: 4D radar) and detailed performance for each weather condition. Best in **bold**, second in underline.

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)	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 (Ours)	L+4DR	<b>85.0</b>	<b>76.6</b>	<b>69.4</b>	<b>74.4</b>	<b>72.3</b>	<b>65.7</b>	<b>93.4</b>	<b>90.4</b>	<b>83.0</b>
1	PointPillars	L	79.9	72.7	67.0	59.9	55.6	50.5	85.5	78.2	72.0
	InterFusion	L+4DR	66.1	64.0	56.9	74.0	70.6	64.5	91.6	87.4	82.0
	L4DR (Ours)	L+4DR	<b>77.9</b>	<b>73.2</b>	<b>67.8</b>	<b>75.4</b>	<b>72.1</b>	<b>66.7</b>	<b>93.8</b>	<b>91.0</b>	<b>83.2</b>
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	<b>63.2</b>	57.8	52.9	77.3	71.1	66.2
	L4DR (Ours)	L+4DR	<b>68.5</b>	<b>56.4</b>	<b>49.3</b>	63.1	<b>59.9</b>	<b>55.1</b>	<b>82.7</b>	<b>70.8</b>	<b>70.7</b>
3	PointPillars	L	44.5	31.9	27.0	40.2	37.7	34.0	53.2	46.7	41.8
	InterFusion	L+4DR	41.2	33.1	27.0	52.9	49.2	44.8	59.9	57.7	53.1
	L4DR (Ours)	L+4DR	<b>46.2</b>	<b>41.4</b>	<b>34.6</b>	<b>53.5</b>	<b>50.6</b>	<b>46.2</b>	<b>72.2</b>	<b>67.7</b>	<b>60.9</b>
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 (Ours)	L+4DR	<b>26.9</b>	<b>26.2</b>	<b>21.6</b>	<b>33.1</b>	<b>30.7</b>	<b>27.9</b>	<b>30.3</b>	<b>29.7</b>	<b>26.3</b>

Table 2: Quantitative results of different methods on Vod-Fog dataset with the KITTI metric under various fog levels.

KONG, and Wijaya 2022), InterFusion (Wang et al. 2022b) and 3D-LRF (Chae, Kim, and Yoon 2024). The results in Table 1 highlight the superior performance of our L4DR model on the K-Radar dataset. Our L4DR model surpasses 3D-LRF by 8.3% in total  $AP_{3D}$ . This demonstrates that compared to previous fusion, our method utilizes the advantages of LiDAR and 4D Radar more effectively. We observed that all methods perform better in many adverse weather conditions (e.g., Overcast, Fog, etc.) than in normal weather. A possible reason is the distribution differences of labeled objects in different weather conditions. We discuss this counter-intuitive phenomenon in the extended version and other valuable results, such as different IoU thresholds and version labels.

### Results on Vod-Fog Simulated Dataset

We evaluated our L4DR model in comparison with LiDAR and 4D radar fusion methods using the Vod-Fog dataset using the KITTI metrics across varying levels of fog. Table 2 demonstrates that our L4DR model outperforms LiDAR-only PointPillars in different difficulty categories and fog intensities. Particularly in the most severe fog conditions (fog level = 4), our L4DR model achieves performance improve-

ments of 17.43%, 17.8%, and 24.81% mAP in moderate difficulty categories, surpassing the gains obtained by InterFusion. Furthermore, our approach consistently exhibits superior performance compared to InterFusion across various scenarios, showcasing the adaptability of our L4DR fusion under adverse weather conditions.

Methods	Modality	Entire Area			Driving Area		
		Car	Ped.	Cyc.	Car	Ped.	Cyc.
Pointpillars	4DR	39.7	31.0	65.1	71.6	40.5	87.8
LXL	4DR	32.8	39.7	68.1	70.3	47.3	87.9
FUTR3D	C+4DR	46.0	35.1	66.0	78.7	43.1	86.2
BEVFusion	C+4DR	37.9	41.0	69.0	70.2	45.9	89.5
RCFusion	C+4DR	41.7	39.0	68.3	71.9	47.5	88.3
LXL	C+4DR	42.3	49.5	77.1	72.2	58.3	88.3
Pointpillars	L	66.0	55.6	75.0	88.7	68.4	88.4
InterFusion	L+4DR	66.5	64.5	78.5	90.7	72.0	88.7
L4DR (Ours)	L+4DR	<b>69.1</b>	<b>66.2</b>	<b>82.8</b>	<b>90.8</b>	<b>76.1</b>	<b>95.5</b>

Table 3: Comparing L4DR with different state-of-the-art methods of different modalities on VoD dataset with the VoD metrics. The best performances are highlighted in **bold**.

Module				3D mAP				
MME	FAD	{IM} <sup>2</sup>	MSGF	W/o Fog	$\mathcal{L} = 1$	$\mathcal{L} = 2$	$\mathcal{L} = 3$	$\mathcal{L} = 4$
				70.3	68.9	53.8	38.8	8.92
✓				77.1	75.4	63.2	52.3	23.4
✓	✓			78.7	77.6	63.3	52.2	24.7
✓	✓	✓		78.1	76.8	62.0	52.3	26.5
✓	✓	✓	✓	<b>79.8</b>	<b>78.8</b>	<b>64.7</b>	<b>53.3</b>	<b>28.9</b>

Table 4: Effect of each component proposed in L4DR, tested on Vod-Fog split. 3D mAP represents the average 3D AP of all classes in Mod. difficulty.  $\mathcal{L}$  denotes fog level.

## Results on VoD Dataset

We compared our L4DR fusion performance with different state-of-the-art methods of different modalities on the VoD dataset with *VoD metric*. As shown in Table 3, our L4DR fusion outperforms the existing LiDAR and 4D radar fusion method InterFusion (Wang et al. 2022b) in all categories. We outperformed by 6.8% in the Cyc. class in the Driving Area. Meanwhile, our L4DR also significantly outperforms other modality-based state-of-the-art methods such as LXL (Xiong et al. 2024). These experimental results demonstrate that our method can comprehensively fuse the two modalities of LiDAR and 4D radar. As a consequence, our L4DR method also shows superior performance even in clear weather.

## Ablation study

**Effect of each component.** We systematically evaluated each component, with the results summarized in Table 4. The 1<sup>st</sup> row represents the performance of the LiDAR-only baseline model. Subsequent 2<sup>nd</sup> row and 3<sup>rd</sup> row are fused by directly concatenating the BEV features from LiDAR and 4D radar modality. The results showing the enhancements were observed with the addition of MME and FAD respectively, highlighting that our fusion method fully utilizes the weather robustness of the 4D radar while excellently handling the noise problem of the 4D radar. The 4<sup>th</sup> row indicates that the performance boost from incorporating the {IM}<sup>2</sup> model alone was not substantial, primarily due to feature redundancy introduced by the {IM}<sup>2</sup> backbone. This issue was effectively addressed by utilizing the MSGF module in the 5<sup>th</sup> row, leading to the most optimal performance.

**Comparison with other feature fusion.** We compared different multi-modal feature fusion blocks, including basic concatenation (Concat.) and various attention-based meth-

Fusion	3D mAP				
	W/o Fog	$\mathcal{L} = 1$	$\mathcal{L} = 2$	$\mathcal{L} = 3$	$\mathcal{L} = 4$
Concat.	77.9	76.3	61.9	49.3	17.7
Cross-Attn.	77.2	76.0	63.0	52.7	<b>30.1</b>
Self-Attn.	78.4	77.4	64.3	52.8	25.8
SE Block	77.3	77.9	63.8	50.1	25.0
CBAM Block	78.0	78.1	64.0	52.3	26.4
MSGF (Ours)	<b>79.8</b>	<b>78.8</b>	<b>64.7</b>	<b>53.2</b>	28.8

Table 5: Comparison of multi-modal feature fusion methods on Vod-Fog val split. 3D mAP represents the average 3D AP of all classes in Mod. difficulty.  $\mathcal{L}$  denotes fog level.

$\lambda$ in testing	Denoise Rate	Recall	mIoU	PA
0.1	88.76	84.35	38.07	88.43
0.2	94.68	78.04	50.11	93.45
0.3	96.51	73.15	54.14	94.78
0.5	97.92	65.52	54.39	95.52

Table 6: Denoise rate, recall, mIoU, and PA (Point Accuracy) of foreground segmentation in FAD ( $\lambda$  in training is 0.3).

ods such as Transformer-based (Vaswani et al. 2017) Cross-Modal Attention (Cross-Attn.) and Self-Attention (Self-Attn.), SE Block (Hu, Shen, and Sun 2018), and CBAM Block (Woo et al. 2018), see extended version for detailed fusion implements. Experimental results (Table 5) show that while attention mechanisms outperform concatenation to some extent, they do not effectively address the challenge of fluctuating features under varying weather conditions. In contrast, our proposed MSGF method, focusing on significant features of LiDAR and 4D radar, achieves superior performance and robustness under different weather.

**Foreground semantic segmentation results in FAD** To explore the 4D radar denoising ability of FAD, we tested different  $\lambda$  in the inference stage, as shown in Table 6. With hyperparameter  $\lambda$  increase, recall gradually decreases, and denoise rate, mIoU, and PA gradually increase. We find that the performance of the 3D object detection is not excellent when  $\lambda$  is relatively high (0.5). The reason is missing foreground points is more unacceptable than retaining some noise for object detection. Additionally, we also discuss the effect of  $\lambda$  on the detection performance in the extended version.

## Conclusion

In this paper, we analyzed the challenges of fusing LiDAR and 4D radar in adverse weather and proposed L4DR, an effective LiDAR and 4D radar fusion method. We provide an innovative and feasible solution for achieving weather-robust outdoor 3D object detection in various weather conditions. Our experiments on VoD and K-Radar datasets have demonstrated the effectiveness and superiority of our method in various simulated fog levels and real-world adverse weather. In summary, our proposed L4DR paves the way for enhanced reliability in autonomous driving and other applications.

**Limitations.** While the introduction of the {IM}<sup>2</sup> and MSGF modules has allowed the model to focus on more salient features, it inevitably introduces additional computations that reduce the computational efficiency to a certain extent. The inference speed is reduced to about **10 FPS**, which is just enough to satisfy the real-time threshold (equal to the LiDAR acquisition frequency), but the computational performance optimization is a valuable future area of research.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No.62171393), and the Fundamental Research Funds for the Central Universities (No.20720220064).

## References

- Balal, N.; Pinhasi, G.; and Pinhasi, Y. 2016. Atmospheric and Fog Effects on Ultra-Wide Band Radar Operating at Extremely High Frequencies. *Sensors*, 16(5): 751.
- Bijelic, M.; Gruber, T.; Mannan, F.; Kraus, F.; Ritter, W.; Dietmayer, K.; and Heide, F. 2020. Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather. In *CVPR*.
- 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 (CVPR)*, 15162–15172.
- Charron, N.; Phillips, S.; and Waslander, S. L. 2018. De-Noising of Lidar Point Clouds Corrupted by Snowfall. In *CRV*.
- Deng, J.; Shi, S.; Li, P.; Zhou, W.; Zhang, Y.; and Li, H. 2021. Voxel R-CNN: Towards High Performance Voxel-based 3D Object Detection. *AAAI*, 35.
- Deng, J.; Ye, W.; Wu, H.; Huang, X.; Xia, Q.; Li, X.; Fang, J.; Li, W.; Wen, C.; and Wang, C. 2025. CMD: A Cross Mechanism Domain Adaptation Dataset for 3D Object Detection. In *European Conference on Computer Vision*, 219–236. Springer.
- Ghita, A.; Antoniusen, B.; Zimmer, W.; Greer, R.; Creß, C.; Møgelmoose, A.; Trivedi, M. M.; and Knoll, A. C. 2024. ActiveAnno3D—An Active Learning Framework for Multi-Modal 3D Object Detection. *arXiv preprint arXiv:2402.03235*.
- Golovachev, Y.; Etinger, A.; Pinhasi, G. A.; and Pinhasi, Y. 2018. Millimeter wave high resolution radar accuracy in fog conditions—theory and experimental verification. *Sensors*, 18(7): 2148.
- Hahner, M.; Sakaridis, C.; Bijelic, M.; Heide, F.; Yu, F.; Dai, D.; and Van Gool, L. 2022. LiDAR Snowfall Simulation for Robust 3D Object Detection. In *CVPR*.
- 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 *ICCV*.
- Han, Z.; Wang, J.; Xu, Z.; Yang, S.; He, L.; Xu, S.; and Wang, J. 2023. 4D Millimeter-Wave Radar in Autonomous Driving: A Survey. *arXiv 2023. arXiv preprint arXiv:2306.04242*.
- Heinzler, R.; Piewak, F.; Schindler, P.; and Stork, W. 2020. CNN-Based Lidar Point Cloud De-Noising in Adverse Weather. *IEEE Robotics and Automation Letters*, 5.
- Hosseinpour, H.; Samadzadegan, F.; and Javan, F. D. 2022. CMGFNet: A deep cross-modal gated fusion network for building extraction from very high-resolution remote sensing images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 184: 96–115.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7132–7141.
- Huang, X.; Wu, H.; Li, X.; Fan, X.; Wen, C.; and Wang, C. 2024a. Sunshine to rainstorm: Cross-weather knowledge distillation for robust 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2409–2416.
- Huang, X.; Xu, Z.; Wu, H.; Wang, J.; Xia, Q.; Xia, Y.; Li, J.; Gao, K.; Wen, C.; and Wang, C. 2024b. L4DR: LiDAR-4DRadar Fusion for Weather-Robust 3D Object Detection. *arXiv:2408.03677*.
- Kilic, V.; Hegde, D.; Sindagi, V. A.; Cooper, A.; Foster, M.; and Patel, V. M. 2021. Lidar Light Scattering Augmentation (LISA): Physics-based Simulation of Adverse Weather Conditions for 3D Object Detection. *ArXiv*.
- 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 *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12689–12697.
- Li, G.; Qian, M.; and Xia, G.-S. 2024. Unleashing Unlabeled Data: A Paradigm for Cross-View Geo-Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16719–16729.
- Li, Y.-J.; O’Toole, M.; and Kitani, K. 2023. St-mvdnet++: Improve vehicle detection with lidar-radar geometrical augmentation via self-training. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Li, Y.-J.; Park, J.; O’Toole, M.; and Kitani, K. 2022. Modality-Agnostic Learning for Radar-Lidar Fusion in Vehicle Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 918–927.
- 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. 2022. K-Radar: 4D Radar Object Detection for Autonomous Driving in Various Weather Conditions. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 3819–3829. Curran Associates, Inc.
- Palfy, 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, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. *arXiv:1706.02413*.
- Qian, K.; Zhu, S.; Zhang, X.; and Li, L. E. 2021. Robust Multimodal Vehicle Detection in Foggy Weather Using Complementary Lidar and Radar Signals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 444–453.
- Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020. PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection. In *CVPR*.
- Shi, S.; Jiang, L.; Deng, J.; Wang, Z.; Guo, C.; Shi, J.; Wang, X.; and Li, H. 2022. PV-RCNN++: Point-Voxel Feature Set Abstraction With Local Vector Representation for 3D Object Detection. *Int. J. Comput. Vision*, 131.

- Song, J.; Zhao, L.; and Skinner, K. A. 2024a. LiRaFusion: Deep Adaptive LiDAR-Radar Fusion for 3D Object Detection. *arXiv preprint arXiv:2402.11735*.
- Song, J.; Zhao, L.; and Skinner, K. A. 2024b. LiRaFusion: Deep Adaptive LiDAR-Radar Fusion for 3D Object Detection. *arXiv preprint arXiv:2402.11735*.
- Sun, S.; and Zhang, Y. D. 2021. 4D automotive radar sensing for autonomous vehicles: A sparsity-oriented approach. *IEEE Journal of Selected Topics in Signal Processing*, 15(4): 879–891.
- Team, O.; et al. 2020. Openpcdet: An open-source toolbox for 3d object detection from point clouds.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, H.; Shi, C.; Shi, S.; Lei, M.; Wang, S.; He, D.; Schiele, B.; and Wang, L. 2023a. DSVT: Dynamic Sparse Voxel Transformer With Rotated Sets. In *CVPR*.
- 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*.
- Wang, L.; Zhang, X.; Xv, B.; Zhang, J.; Fu, R.; Wang, X.; Zhu, L.; Ren, H.; Lu, P.; Li, J.; and Liu, H. 2022b. InterFusion: Interaction-based 4D Radar and LiDAR Fusion for 3D Object Detection. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 12247–12253.
- Wang, Y.; Deng, J.; Li, Y.; Hu, J.; Liu, C.; Zhang, Y.; Ji, J.; Ouyang, W.; and Zhang, Y. 2023b. Bi-LRFusion: Bi-Directional LiDAR-Radar Fusion for 3D Dynamic Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13394–13403.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, 3–19.
- Wu, H.; Wen, C.; Shi, S.; Li, X.; and Wang, C. 2023. Virtual Sparse Convolution for Multimodal 3D Object Detection. In *CVPR*.
- Wu, H.; Zhao, S.; Huang, X.; Wen, C.; Li, X.; and Wang, C. 2024. Commonsense Prototype for Outdoor Unsupervised 3D Object Detection. *arXiv preprint arXiv:2404.16493*.
- Xia, Q.; Ye, W.; Wu, H.; Zhao, S.; Xing, L.; Huang, X.; Deng, J.; Li, X.; Wen, C.; and Wang, C. 2024. HINTED: Hard Instance Enhanced Detector with Mixed-Density Feature Fusion for Sparsely-Supervised 3D Object Detection. In *CVPR*, 15321–15330.
- Xiong, W.; Liu, J.; Huang, T.; Han, Q.-L.; Xia, Y.; and Zhu, B. 2024. LXL: LiDAR Excluded Lean 3D Object Detection With 4D Imaging Radar and Camera Fusion. *IEEE Transactions on Intelligent Vehicles*, 9(1): 79–92.
- Xu, G.; Khan, A.; Moshayedi, A. J.; Zhang, X.; and Shuxin, Y. 2022. The Object Detection, Perspective and Obstacles In Robotic: A Review. *EAI Endorsed Transactions on AI and Robotics*, 1: 7–15.
- Xu, Q.; Zhou, Y.; Wang, W.; Qi, C. R.; and Anguelov, D. 2021. SPG: Unsupervised Domain Adaptation for 3D Object Detection via Semantic Point Generation. In *ICCV*.
- Yan, Y.; Mao, Y.; and Li, B. 2018. SECOND: Sparsely Embedded Convolutional Detection. *Sensors*, 18.
- Yang, Z.; Sun, Y.; Liu, S.; and Jia, J. 2020. 3DSSD: Point-Based 3D Single Stage Object Detector. In *CVPR*.