

# Bi-LRFusion: Bi-Directional LiDAR-Radar Fusion for 3D Dynamic Object Detection

Yingjie Wang<sup>1,4</sup> Jiajun Deng<sup>2\*</sup> Yao Li<sup>1</sup> Jinshui Hu<sup>3</sup> Cong Liu<sup>3</sup> Yu Zhang<sup>1</sup> Jianmin Ji<sup>1</sup>  
Wanli Ouyang<sup>4</sup> Yanyong Zhang<sup>1\*</sup>

<sup>1</sup>University of Science and Technology of China <sup>2</sup>University of Sydney <sup>3</sup>iFLYTEK <sup>4</sup>Shanghai AI Laboratory

## Abstract

LiDAR and Radar are two complementary sensing approaches in that LiDAR specializes in capturing an object's 3D shape while Radar provides longer detection ranges as well as velocity hints. Though seemingly natural, how to efficiently combine them for improved feature representation is still unclear. The main challenge arises from that Radar data are extremely sparse and lack height information. Therefore, directly integrating Radar features into LiDAR-centric detection networks is not optimal. In this work, we introduce a bi-directional LiDAR-Radar fusion framework, termed Bi-LRFusion, to tackle the challenges and improve 3D detection for dynamic objects. Technically, Bi-LRFusion involves two steps: first, it enriches Radar's local features by learning important details from the LiDAR branch to alleviate the problems caused by the absence of height information and extreme sparsity; second, it combines LiDAR features with the enhanced Radar features in a unified bird's-eye-view representation. We conduct extensive experiments on nuScenes and ORR datasets, and show that our Bi-LRFusion achieves state-of-the-art performance for detecting dynamic objects. Notably, Radar data in these two datasets have different formats, which demonstrates the generalizability of our method. Codes are available at <https://github.com/JessieW0806/Bi-LRFusion>.

## 1. Introduction

LiDAR has been considered as the primary sensor in the perception subsystem of most autonomous vehicles (AVs) due to its capability of providing accurate position measurements [9, 16, 32]. However, in addition to object positions, AVs are also in an urgent need for estimating the motion state information (e.g., velocity), especially for dynamic objects. Such information cannot be measured by

\*Corresponding Author: Jiajun Deng and Yanyong Zhang.

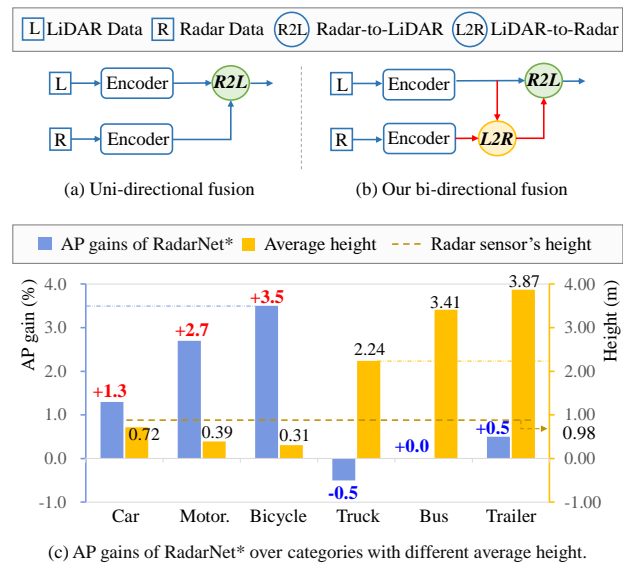


Figure 1. An illustration of (a) uni-directional LiDAR-Radar fusion mechanism, (b) our proposed bi-directional LiDAR-Radar fusion mechanism, and (c) the average precision gain (%) of uni-directional fusion method RadarNet\* against the LiDAR-centric baseline CenterPoint [40] over categories with different average height (m). We use \* to indicate it is re-produced by us on the CenterPoint. The improvement by involving Radar data is not consistent for objects with different height, i.e., taller objects like truck, bus and trailer do not enjoy as much performance gain. Note that all height values are transformed to the LiDAR coordinate system.

LiDAR sensors since they are insensitive to motion. As a result, millimeter-wave Radar (referred to as Radar in this paper) sensors are engaged because they are able to infer the object's relative radial velocity [21] based upon the Doppler effect [28]. Besides, on-vehicle Radar usually offers longer detection range than LiDAR [36], which is particularly useful on highways and expressways. In the exploration of combining LiDAR and Radar data for ameliorating 3D dynamic object detection, the existing approaches [22, 25, 36] follow the common mechanism of *uni-directional* fusion,

as shown in Figure 1 (a). Specifically, these approaches directly utilize the Radar data/feature to enhance the LiDAR-centric detection network without first improving the quality of the feature representation of the former.

However, independently extracted Radar features are not enough for refining LiDAR features, since Radar data are extremely sparse and lack the height information<sup>1</sup>. Specifically, taking the data from the nuScenes dataset [4] as an example, the 32-beam LiDAR sensor produces approximately 30,000 points, while the Radar sensor only captures about 200 points for the same scene. The resulting Radar bird’s eye view (BEV) feature hardly attains valid local information after being processed by local operators (*e.g.*, the neighbors are most likely empty when a non-empty Radar BEV pixel is convolved by convolutional kernels). Besides, on-vehicle Radar antennas are commonly arranged horizontally, hence missing the height information in the vertical direction. In previous works, the height values of the Radar points are simply set as the ego Radar sensor’s height. Therefore, when features from Radar are used for enhancing the feature of LiDAR, the problematic height information of Radar leads to unstable improvements for objects with different heights. For example, Figure 1 (c) illustrates this problem. The representative method RadarNet falls short in the detection performance for tall objects – the truck class even experiences 0.5% AP degradation after fusing the Radar data.

In order to better harvest the benefit of LiDAR and Radar fusion, our viewpoint is that Radar features need to be more powerful before being fused. Therefore, we first enrich the Radar features – with the help of LiDAR data – and then integrate the enriched Radar features into the LiDAR processing branch for more effective fusion. As depicted in Figure 1 (b), we refer to this scheme as *bi-directional* fusion. And in this work, we introduce a framework, *Bi-LRFusion*, to achieve this goal. Specifically, Bi-LRFusion first encodes BEV features for each modality individually. Next, it engages the query-based LiDAR-to-Radar (L2R) height feature fusion and query-based L2R BEV feature fusion, in which we query and group LiDAR points and LiDAR BEV features that are close to the location of each non-empty grid cell on the Radar feature map, respectively. The grouped LiDAR raw points are aggregated to formulate pseudo-Radar height features, and the grouped LiDAR BEV features are aggregated to produce pseudo-Radar BEV features. The generated pseudo-Radar height and BEV features are fused to the Radar BEV features through concatenation. After enriching the Radar features, Bi-LRFusion then performs the Radar-to-LiDAR (R2L) fusion in a unified BEV representation. Finally, a BEV detection network

<sup>1</sup>This shortcoming is due to today’s Radar technology, which may likely change as the technology advances very rapidly, *e.g.*, new-generation 4D Radar sensors [3].

consisting of a BEV backbone network and a detection head is applied to output 3D object detection results.

We validate the merits of bi-directional LiDAR-Radar fusion via evaluating our Bi-LRFusion on nuScenes and Oxford Radar RobotCar (ORR) [1] datasets. On nuScenes dataset, Bi-LRFusion improves the mAP(↑) by 2.7% and reduces the mAVE(↓) by 5.3% against the LiDAR-centric baseline CenterPoint [40], and remarkably outperforms the strongest counterpart, *i.e.*, RadarNet, in terms of AP by absolutely 2.0% for cars and 6.3% for motorcycles. Moreover, Bi-LRFusion generalizes well on the ORR dataset, which has a different Radar data format and achieves 1.3% AP improvements for vehicle detection.

In summary, we make the following contributions:

- We propose a bi-directional fusion framework, namely Bi-LRFusion, to combine LiDAR and Radar features for improving 3D dynamic object detection.
- We devise the query-based L2R height feature fusion and query-based L2R BEV feature fusion to enrich Radar features with the help of LiDAR data.
- We conduct extensive experiments to validate the merits of our method and show considerably improved results on two different datasets.

## 2. Related Work

### 2.1. 3D Object Detection with Only LiDAR Data

In the literature, 3D object detectors built on LiDAR point clouds are categorized into two directions, *i.e.*, point-based methods and voxel-based methods.

**Point-based Methods.** Point-based methods maintain the precise position measurement of raw points, extracting point features by point networks [23, 24] or graph networks [33, 34]. The early work PointRCNN [27] follows the two-stage pipeline to first produce 3D proposals from pre-assigned anchor boxes over sampled key points, and then refines these coarse proposals with region-wise features. STD [39] proposes the sparse-to-dense strategy for better proposal refinement. A more recent work 3DSSD [38] further introduces feature based key point sampling strategy as a complement of previous distance based one, and develops a one-stage object detector operating on raw points.

**Voxel-based Methods.** Voxel-based methods first divide points into regular grids, and then leverage the convolutional neural networks [8, 9, 15, 26, 35, 41] and Transformers [11, 20] for feature extraction and bounding box prediction. SECOND [35] reduces the computational overhead of dense 3D CNNs by applying sparse convolution. PointPillars [15] introduces a pillar representation (a particular form of the voxel) to formulate with 2D convolutions. To simplify and improve previous 3D detection pipelines, CenterPoint [40] designs an anchor-free one-stage detector, which

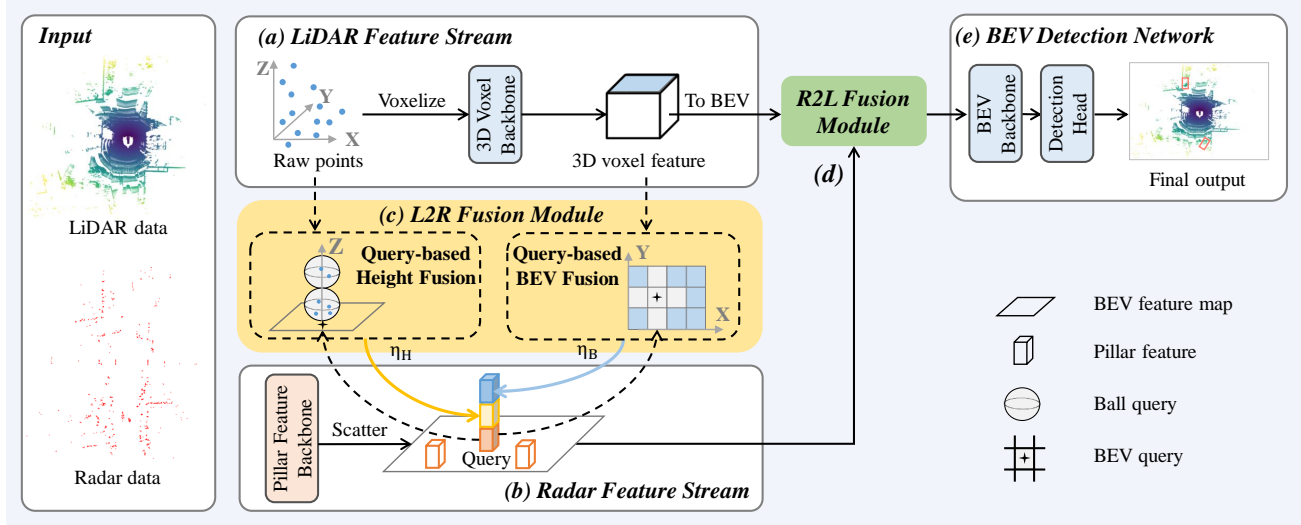


Figure 2. An overview of the proposed Bi-LRFusion framework. Bi-LRFusion includes five main components: (a) a LiDAR feature stream to encode LiDAR BEV features from LiDAR data, (b) a Radar feature stream to encode Radar BEV features from Radar data, (c) a LiDAR-to-Radar (L2R) fusion module composed of a query-based height feature fusion block and a query-based BEV feature fusion block, in which we enhance Radar features from LiDAR raw points and LiDAR features, (d) a Radar-to-LiDAR (R2L) fusion module to fuse back the enhanced Radar features to the LiDAR-centric detection network, and (e) a BEV detection network that uses the features from the R2L fusion module to predict 3D bounding boxes for dynamic objects.

extracts BEV features from voxelized point clouds to find object centers and regress to 3D bounding boxes. Furthermore, CenterPoint introduces a velocity head to predict the object’s velocity between consecutive frames. In this work, we exploit CenterPoint as our LiDAR-only baseline.

## 2.2. 3D Object Detection with Radar Fusion

Providing larger detection range and additional velocity hints, Radar data show great potential in 3D object detection. However, since they are too sparse to be solely used [4], Radar data are generally explored as the complement of RGB images or LiDAR point clouds. The approaches that fuse Radar data for improving 3D object detectors can be summarized into two categories: one is input-level fusion, and the other is feature-level fusion.

**Input-level Radar Fusion.** RVF-Net [22] develops an early fusion scheme to treat raw Radar points as additional input besides LiDAR point clouds, ignoring the differences in data property between Radar and LiDAR. As studied in [7], these input-level fusion methods directly incorporate Radar raw information into LiDAR branch, which is sensitive to even slight changes of the input data and is also unable to fully utilize the multi-modal feature.

**Feature-level Radar Fusion.** The early work GRIF [14] proposes to extract region-wise features from Radar and camera branches, and to combine them together for robust 3D detection. Recent works generally transform image features to the BEV plane for feature fusion [12, 13, 17]. MVD-Net [25] encodes Radar points’ intensity into a BEV fea-

ture map, and fuses it with LiDAR features to facilitate vehicle detection under foggy weather. The representative work RadarNet [36] first extracts LiDAR and Radar features via modality-specific branches, and then fuses them on the shared BEV perspective. The existing feature-level Radar fusion methods commonly ignore the problems caused by the height missing and extreme sparsity of Radar data, as well as overlooking the information intensity gap when fusing the multi-modal features.

Despite both input-level and feature-level methods have improved 3D dynamic object detection via Radar fusion, they follow the uni-directional fusion scheme. On the contrary, our Bi-LRFusion, for the first time, treats LiDAR-Radar fusion in a bi-directional way. We enhance the Radar feature with the help of LiDAR data to alleviate issues caused by the absence of height information and extreme sparsity, and then fuse it to the LiDAR-centric network to achieve further performance boosting.

## 3. Methodology

In this work, we present Bi-LRFusion, a bi-directional LiDAR-Radar fusion framework for 3D dynamic object detection. As illustrated in Figure 2, the input LiDAR and Radar points are fed into the sibling LiDAR feature stream and Radar feature stream to produce their BEV features. Next, we involve a LiDAR-to-Radar (L2R) fusion step to enhance the extremely sparse Radar features that lack discriminative details. Specifically, for each valid (*i.e.*, non-empty) grid cell on the Radar feature map, we query and

group the nearby LiDAR data (including both raw points and BEV features) to obtain more detailed Radar features. Here, we focus on the height information that are completely missing in the Radar data and the local BEV features that are scarce in the Radar data. Through two query-based feature fusion blocks, we can transfer the knowledge from LiDAR data to Radar features, leading to much-enriched Radar features. Subsequently, we perform Radar-to-LiDAR (R2L) fusion by integrating the enriched Radar features to the LiDAR features in a unified BEV representation. Finally, a BEV detection network composed of a BEV backbone network and a detection head outputs 3D detection results. Below we describe these steps one by one.

### 3.1. Modality-Specific Feature Encoding

**LiDAR Feature Encoding.** LiDAR feature encoding consists of the following steps. Firstly, we divide LiDAR points into 3D regular grids and encode the feature of each grid by a multi-layer perception (MLP) followed by max-pooling. Here, a grid is known as a voxel in the discrete 3D space, and the encoded feature is the voxel feature. After obtaining voxel features, we follow the common practice to exploit a 3D voxel backbone network composed of 3D sparse convolutional layers and 3D sub-manifold convolutional layers [35] for LiDAR feature extraction. Then, the output feature volume is stacked along the Z-axis, producing a LiDAR BEV feature map  $M_L \in \mathbb{R}^{C_1 \times H \times W}$ , where  $(H, W)$  indicate the height and width.

**Radar Feature Encoding.** A Radar point contains the 2D coordinate  $(x, y)$ , the Radar cross-section  $r_{cs}$ , and the timestamp  $t$ . In addition, we may also have<sup>2</sup>the radial velocity of the object in  $X - Y$  directions  $(v_X, v_Y)$ , the dynamic property  $dynProp$ , the cluster validity state  $invalid\_state$ , and the false alarm probability  $pdh0$ . Please note that the Radar point’s value on the Z-axis is set to be the Radar sensor’s height by default. Here, we exploit the pillar [15] representation to encode Radar features, which directly converts the Radar input to a pseudo image in the bird’s eye view (BEV). Then we extract Radar features with a pillar feature network, obtaining a Radar BEV feature map  $M_R \in \mathbb{R}^{C_2 \times H \times W}$ .

### 3.2. LiDAR-to-Radar Fusion

LiDAR-to-Radar (L2R) fusion is the core step in our proposed Bi-LRFusion. It involves two L2R feature fusion blocks, in which we generate pseudo height features, as well as the pseudo local BEV features, by suitably querying the LiDAR features. These pseudo features are then fused into the Radar features to enhance their quality.

**Query-based L2R Height Feature Fusion.** As illustrated

<sup>2</sup>Due to the different 2D millimeter wave Radar sensors installed, the Radar data in the ORR dataset are stored in the form of intensity images and therefore do not include the other features mentioned here.

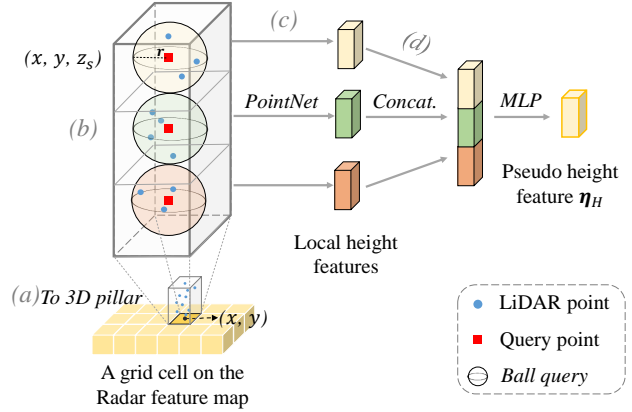


Figure 3. An illustration of query-based L2R height feature fusion (QHF) block. QHF involves the following steps: (a) lifting the non-empty grid cell on the Radar feature map to a pillar, and equally dividing the pillar into segments of different heights, (b) querying and grouping neighboring LiDAR points based on the location of each segment’s center, (c) aggregating grouped LiDAR points to get the local height feature of each segment, and (d) merging the segments’ feature together to produce pseudo height feature  $\eta_H$  for the corresponding Radar grid cell.

in Figure 1 (c), directly fusing the Radar features to LiDAR-based 3D detection networks may lead to unsatisfactory results, especially for objects that are taller than the Radar sensor. This is caused by the fact that Radar points do not contain the height measurements. To address this insufficiency, we design the query-based L2R height feature fusion (QHF) block to transfer the height distributions from LiDAR raw points to Radar features.

The pipeline of QHF is depicted in Figure 3. The core innovation of QHF is the height feature querying mechanism. Given the  $l$ -th valid grid cell on the Radar feature map centered at  $(x, y)$ , we first “lift” it from the BEV plane to the 3D space, which results in a pillar of height  $h$ . Then, we evenly divide the pillar into  $M$  segments along the height and assign a query point at each segment’s center. Specifically, let us denote the grid size of the Radar BEV feature map as  $r \times r$ . In order to query LiDAR points without overlap, the radius of the ball query is set to  $r/2$ , and the number of segments  $M$  is set to  $h/2r$ . For the query point in  $s$ -th segment shown in Figure 3, the  $(x, y)$  coordinates are from the central point of the given grid cell, which can be calculated with the grid indices, and grid sizes, together with the boundaries of the Radar points. Further, the height value  $z_s$  of the query point is calculated as:

$$z_s = z_M + r \times (2s - 1), \quad (1)$$

where  $z_M$  is the minimum height value among all the LiDAR points. After establishing the query points, we then apply ball query [24] followed by a PointNet module [23] to aggregate the local height feature  $F_s$  from the grouped

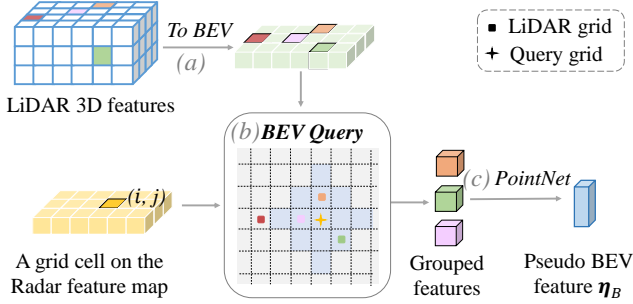


Figure 4. An illustration of query-based L2R BEV feature fusion (QBF) block with the following steps: (a) collapsing LiDAR 3D features to the BEV grids, (b) querying and grouping LiDAR grids that are close to the Radar query grid based on their indices on the BEV, and (c) aggregating the features of the grouped LiDAR grids to generate the pseudo BEV feature  $\eta_B$  for the corresponding Radar grid cell.

LiDAR points. The calculation can be formulated as:

$$\mathbf{F}_s = \max_{k=1,2,\dots,K} \left\{ \Psi(n_s^k) \right\}, \quad (2)$$

where  $n_s^k$  denotes the  $k$ -th grouped LiDAR point in the  $s$ -th ball query segment,  $K$  is the number of grouped points from ball query,  $\Psi(\cdot)$  indicates an MLP, and  $\max(\cdot)$  is the max pooling operation. After obtaining the local height feature of each segment, we concatenate them together and feed the concatenated feature into an MLP to make the output channels have the same dimensions as the Radar BEV feature map. Finally, the output pseudo height feature  $\eta_H^l$  for the  $l$ -th grid cell on the Radar feature map produced by QHF is computed as:

$$\eta_H^l = \text{MLP} \left( \text{Concat} \left( \{ \mathbf{F}_s \}_{s=0}^M \right) \right). \quad (3)$$

**Query-based L2R BEV Feature Fusion.** When extremely sparse Radar BEV pixels are convoluted by convolutional kernels, the resulting Radar BEV feature barely retains valid local information since most of the neighboring pixels are empty. To alleviate this problem, we design a query-based L2R BEV feature fusion block (QBF) that can generate the pseudo BEV feature which is more detailed than the original Radar BEV feature, by querying and grouping the corresponding fine-grained LiDAR BEV features.

The pipeline of QBF is depicted in Figure 4. The core innovation of QBF is the local BEV feature querying mechanism that we describe below. We first collapse LiDAR 3D features to the BEV grid plane [15], forming a set of non-empty LiDAR grids (the same grid size as the Radar grid cell in Figure 4). Given the  $l$ -th non-empty grid cell on the Radar feature map with indices  $(i, j)$ , we query and group LiDAR grid features that are close to the Radar grid cell on the BEV plane. Specifically, towards this goal, we propose a local BEV feature query inspired by Voxel R-CNN [8],

which finds all LiDAR grid features that are within a certain distance from the querying grid based on their grid indices on the BEV plane. Specifically, we exploit the Manhattan distance metric and sample up to  $K$  non-empty LiDAR grids within a specific distance threshold on the BEV plane. The Manhattan distance  $D(\alpha, \beta)$  between indices of LiDAR grids  $\alpha = \{i_\alpha, j_\alpha\}$  and  $\beta = \{i_\beta, j_\beta\}$  can be calculated as:

$$D(\alpha, \beta) = |i_\alpha - i_\beta| + |j_\alpha - j_\beta|, \quad (4)$$

where  $i$  and  $j$  are the indices of the LiDAR grid along the  $X$  and  $Y$  axis. After the  $l$ -th BEV feature query, we obtain the corresponding features of the grouped LiDAR grids. Finally, we apply a PointNet module [23] to aggregate the pseudo Radar BEV feature  $\eta_B^l$ , which can be summarized as:

$$\eta_B^l = \max_{k=1,2,\dots,K} \left\{ \Psi'(F_l^k) \right\}, \quad (5)$$

where  $F_l^k$  denotes the  $k$ -th grouped LiDAR grid feature in the  $l$ -th BEV query mechanism.

### 3.3. Radar-to-LiDAR Fusion

After enriching the Radar BEV features with the pseudo height feature  $\eta_H$  and pseudo BEV features  $\eta_B$ , we obtain the enhanced Radar BEV features that have 96 channels.

In this step, we fuse the enhanced Radar BEV features to the LiDAR-based 3D detection pipeline so as to incorporate valuable clues such as velocity information. Specifically, we concatenate the two BEV features in the channel-wise fashion following the practice from [19, 36]. Before forwarding to the BEV detection network, we also apply a convolution-based BEV encoder to help curb the effect of misalignment between multi-modal BEV features. The BEV encoder adjusts the fused BEV feature to 512 through three 2D convolution blocks.

### 3.4. BEV Detection Network

Finally, the combined LiDAR and Radar BEV features are fed to the BEV detection network to output the results. The BEV detection network consists of a BEV network and a detection head. The BEV network is composed of several 2D convolution blocks, which generate center features forwarding to the detection head. As such, we use a class-specific center heatmap head to predict the center location of all dynamic objects and a few regression heads to estimate the object size, rotation, and velocity based on the center features. We combine all heatmap and regression losses in one common objective and jointly optimize them following the baseline CenterPoint [40].

Table 1. Comparison with other methods on the nuScenes validation set. We add “\*” to indicate that it is a reproducing version based on CenterPoint [40]. We group the dynamic targets to (1) similar-height objects and (2) tall objects according to the Radar sensor’s height.

Method	Modality	mAVE ↓	mAP ↑	AP ↑ of Group 1				AP ↑ of Group 2		
				Car	Motor.	Bicycle	Ped.	Truck	Bus	Trailer
PointPillars [15]	L	34.2	47.1	80.7	26.7	5.3	70.8	49.4	62.0	34.9
SECOND [35]		32.1	52.9	81.7	40.1	18.2	77.4	51.9	66.0	38.2
CenterPoint [40]		30.3	59.3	84.6	54.7	34.9	83.9	54.4	66.7	36.7
RVF-Net [22]	L + R	-	54.9	-	-	-	-	-	-	-
RadarNet [36]		-	-	84.5	52.9	-	-	-	-	-
RadarNet* [36]		26.2	60.4	85.9	57.4	38.4	84.1	53.9	66.7	37.2
Bi-LRFusion		<b>25.0</b>	<b>62.0</b>	<b>86.5</b>	<b>59.2</b>	<b>42.0</b>	<b>84.4</b>	<b>55.2</b>	<b>67.9</b>	<b>38.4</b>

## 4. Experiments

We evaluate Bi-LRFusion on both the nuScenes and the Oxford Radar RobotCar (ORR) datasets and conduct ablation studies to verify our proposed fusion modules. We further show the advantages of Bi-LRFusion on objects with different heights/velocities.

### 4.1. Datasets and Evaluation Metrics

**NuScenes Dataset.** NuScenes [4] is a large-scale dataset for 3D detection including camera images, LiDAR points, and Radar points. We mainly adopt two metrics in our evaluation. The first one is the average precision (AP) with a match threshold of 2D center distance on the ground plane. The second one is AVE which stands for absolute velocity error in  $m/s$  – its decrease represents more accurate velocity estimation. We average the two metrics over all 7 dynamic classes (mAP, mAVE), following the official evaluation.

**ORR Dataset.** The ORR dataset, mainly for localization and mapping tasks, is a challenging dataset including camera images, LiDAR points, Radar scans, GPS and INS ground truth. We split the first data record into 7,064 frames for training and 1,760 frames for validation. As this dataset does not provide ground truth of object annotation, we therefore follow the MVDet [25] to generate 3D boxes of vehicles. For evaluation metrics, we use AP of oriented bounding boxes in BEV to validate the vehicle detection performance via the CoCo evaluation framework [18].

### 4.2. Implementation Details

**LiDAR Input.** As allowed in the nuScenes submission rule, we accumulate 10 LiDAR sweeps to form a denser point cloud. we set detection range to  $(-54.0, 54.0)m$  for the  $X, Y$  axis, and  $(-5.0, 3.0)m$  for  $Z$  axis with a voxel size of  $(0.075, 0.075, 0.2)m$ . For ORR dataset, we set the point cloud range to  $(-69.12, 69.12)m$  for  $X$  and  $Y$  axis,  $(-5.0, 2.0)m$  for  $Z$  axis and voxel size to  $(0.32, 0.32, 7.0)m$ .

**Radar Input in NuScenes.** Radar data are collected from 5 long-range Radar sensors and stored as BIN files, which is the same format as LiDAR point cloud. We stack points captured by five Radar sensors into full-view Radar point

clouds. We also accumulate 6 sweeps to form denser Radar points. The detection range of Radar data is consistent with the LiDAR range. The voxel size is  $(0.6, 0.6, 8.0)m$  for nuScenes and  $(0.32, 0.32, 7.0)m$  for ORR dataset. We transfer the 2D position and velocity of Radar points from the Radar coordinates to the LiDAR coordinates.

**Radar Input in ORR.** The Radar data are collected by a vehicle equipped with only one spinning Millimetre-Wave FMCW Radar sensor and are saved as PNG files. Therefore, we need to convert the Radar images to Radar point clouds. First, we use cen2019 [5] as a feature extractor to extract feature points from Radar images. Attributes of every extracted point include  $x, y, z, rcs, t$ , where  $rcs$  is represented by gray-scale values on the Radar image, and  $z$  is set to 0. To decrease the noise points and ghost points due to multi-path effects, we apply a geometry-probabilistic filter [10] and the number of points is reduced to around 1,000 per frame. Second, since Radar’s considerable scanning delay could cause the lack of synchronization with the LiDAR, we compensate for the ego-motion via SLAM [30]. Please refer to supplementary materials for more details.

### 4.3. Comparison on NuScenes Dataset

We evaluate our Bi-LRFusion with AP (%) and AVE (m/s) of 7 dynamic classes on both validation and test sets.

First, we compare our Bi-LRFusion with a few top-ranked methods on the nuScenes validation set, including the vanilla CenterPoint [40] as our baseline, the LiDAR-Radar fusion models and other LiDAR-only models on nuScenes benchmark. For a fair comparison, we use consistent settings with the original papers and reproduce the results on our own. Table 1 summarizes the results. Overall, our results show solid performance improvements. The SOTA LiDAR-Radar fusion method, RadarNet [36], mainly focused on improving the AP values for cars and motorcycles in their study. Compared to RadarNet, we can further improve the AP by **+2.0%** for cars and **+6.3%** for motorcycles. In addition, when we consider all dynamic object categories, Bi-LRFusion can increase the mAP(↑) by **+2.7%** and improve mAVE(↓) by **-5.3%** compared to CenterPoint. Meanwhile, our approach surpasses the repro-

Table 2. Comparison with other methods on the nuScenes test set. We add “\*” to indicate that we reproduce and submit the result based on CenterPoint [40]. We group the dynamic targets to (1) similar-height objects and (2) tall objects according to the Radar sensor’s height.

Method	Modality	mAVE ↓	mAP ↑	Car	AP ↑ of Group 1			AP ↑ of Group 2		
					Motor.	Bicycle	Ped.	Truck	Bus	Trailer
PointPillars [15]	L	31.6	30.5	68.4	27.4	1.1	59.7	23.0	28.2	23.4
InfoFocus [31]		-	39.5	77.9	29.0	6.1	64.3	31.4	44.8	37.3
PointPillars+ [29]		27.0	40.1	76.0	34.2	14.0	64.0	31.0	32.1	36.6
SSN [42]		26.6	42.6	82.4	48.9	24.6	75.6	41.8	46.1	48.0
CVCNet [6]		26.8	55.3	82.7	59.1	31.3	79.8	46.1	46.6	49.4
CenterPoint [40]		28.8	60.3	85.2	59.5	30.7	84.6	53.5	63.6	56.0
RadarNet* [36]	L + R	27.9	61.9	86.0	59.6	31.9	<b>84.8</b>	53.0	62.8	55.4
Bi-LRFusion		<b>25.7</b>	<b>63.1</b>	<b>87.2</b>	<b>59.9</b>	<b>34.0</b>	84.7	<b>54.8</b>	<b>63.7</b>	<b>57.3</b>

Table 3. Comparison with other methods of vehicle detection accuracy on the ORR dataset.

Method	Modality	AP ↑	
		IoU=0.5	IoU=0.8
PIXOR [37]	L	72.8	41.2
PointPillars [15]		85.7	58.3
DEF [2]	L + R	86.6	46.2
MVDet [25]		90.9	<b>74.6</b>
Bi-LRFusion		<b>92.2</b>	74.4

duced RadarNet\* by +1.6% in mAP and -1.2% in mAVE.

We also compare our method with several top-performing models on the test set. The detailed results are listed in Table 2. The results show that Bi-LRFusion gives the best average results, in terms of both mAP and mAVE. For six out of seven individual object categories, it has the best AP results. Even for the only exception category (the pedestrian), its AP is the second best, with only 0.1% lower than the best. Overall, the results on val and test sets consistently demonstrate the effectiveness of Bi-LRFusion.

Furthermore, we qualitatively compare with CenterPoint and Bi-LRFusion on the nuScenes dataset. Figure 5 shows a few visualized detection results, which demonstrates that the enhanced Radar data from L2R feature fusion module can indeed better eliminate the missed detections.

#### 4.4. Comparison on ORR Dataset

To further validate the effectiveness of our proposed Bi-LRFusion, we conduct experiments on the challenging ORR dataset. Due to the lack of ground truth annotations, we only report the AP for cars with IoU thresholds of 0.5 and 0.8, following the common COCO protocol.

We compare our Bi-LRFusion with several LiDAR-only and LiDAR-Radar fusion methods on the ORR dataset. From Table 3, Bi-LRFusion achieves the best AP results with IoU threshold of 0.5, and second best with IoU threshold of 0.8. For IoU threshold of 0.8, its AP is only 0.2% lower than the best, and much higher than other schemes. Table 3 also shows that Bi-LRFusion is effective in handling different Radar data formats.

Table 4. The effect of each proposed component in Bi-LRFusion. It shows that all components contribute to the overall detection performance.

Method	R2L fusion	L2R fusion		mAP ↑	mAVE ↓
		QHF	QBF		
(a)				59.3	30.3
(b)	✓			60.4 (+1.1)	26.2 (-4.1)
(c)	✓	✓		61.3 (+2.0)	25.1 (-5.2)
(d)	✓		✓	61.9 (+2.6)	25.5 (-4.8)
(e)	✓	✓	✓	<b>62.0 (+2.7)</b>	<b>25.0 (-5.3)</b>

#### 4.5. Effect of Different Bi-LRFusion Modules

To understand how each module in Bi-LRFusion affects the detection performance, we conduct ablation studies on the nuScenes validation set and report its mAP and mAVE of all dynamic classes in Table 4.

*Method (a)* is our LiDAR-only baseline CenterPoint, which achieves mAP of 59.3% and mAVE of 30.3%.

*Method (b)* extends (a) by simply fusing the Radar feature via R2L fusion, which improves mAP by 1.1% and mAVE by -4.1%. This indicates that integrating the Radar feature is effective to improve 3D dynamic detection.

*Method (c)* extends (b) by utilizing the raw LiDAR points to enhance the Radar features via query-based L2R height feature fusion (QHF), which leads to an improvement of 2.0% mAP and -5.2% mAVE.

*Method (d)* extends (b) by taking advantage of the detailed LiDAR features on the BEV plane to enhance Radar features via query-based L2R BEV feature fusion (QBF), improving mAP by 2.6% and mAVE by -4.8%.

*Method (e)* is our Bi-LRFusion. By combining all the components, it achieves a gain of 2.7% for mAP and -5.3% for mAVE compared to CenterPoint. By enhancing the Radars features before fusing them into the detection network, our bi-directional LiDAR-Radar fusion framework can considerably improve the detection of moving objects.

#### 4.6. Effect of Object Parameters

We also evaluate the performance gain of Bi-LRFusion for objects with different heights/velocities compared with the LiDAR-centric baseline CenterPoint.

Table 5. mAP and mAVE results for LiDAR-only CenterPoint [40] and Bi-LRFusion for objects with different velocities.

Method	Velocity (m/s)	mAP $\uparrow$	mAVE $\downarrow$
CenterPoint	0 – 0.5	81.3	7.0
	0.5 – 5.0	44.9	57.1
	5.0 – 10.0	61.5	62.7
	$\geq 10.0$	55.4	83.5
Bi-LRFusion	0 – 0.5	<b>82.6</b> (+1.3)	<b>6.9</b> (-0.1)
	0.5 – 5.0	<b>52.1</b> (+7.2)	<b>45.9</b> (-11.2)
	5.0 – 10.0	<b>68.6</b> (+7.1)	<b>51.4</b> (-11.3)
	$\geq 10.0$	<b>65.6</b> (+10.2)	<b>69.7</b> (-13.8)

(a) cars

Method	Velocity (m/s)	mAP $\uparrow$	mAVE $\downarrow$
CenterPoint	0 – 0.5	52.7	6.2
	0.5 – 5.0	31.0	82.4
	5.0 – 10.0	33.8	115.1
	$\geq 10.0$	54.0	255.5
Bi-LRFusion	0 – 0.5	<b>54.0</b> (+1.3)	<b>6.0</b> (-0.2)
	0.5 – 5.0	<b>37.3</b> (+6.3)	<b>79.1</b> (-3.3)
	5.0 – 10.0	<b>42.5</b> (+8.7)	<b>93.4</b> (-21.7)
	$\geq 10.0$	<b>70.2</b> (+16.2)	<b>154.8</b> (-100.7)

(b) motorcycles

Table 6. The percentage (%) of objects with different velocities in nuScenes dataset, *i.e.*, stationary (0 m/s - 0.5 m/s), low-velocity (0.5 m/s - 5.0 m/s), medium-velocity (5.0 m/s - 10.0 m/s), high-velocity ( $\geq 10.0$  m/s). These groups are divided according to [4].

Class	0 – 0.5	0.5 – 5.0	5.0 – 10.0	$\geq 10.0$
Cars	72.6	11.7	11.8	3.9
Motorcycles	69.8	13.0	12.3	4.9

**Effect of Object Height.** We group the dynamic objects into two groups: (1) regular-height objects including cars, motorcycles, bicycles, pedestrians, and (2) tall objects including trucks, buses, and trailers. Note that since millimeter wave is not sensitive to non-rigid targets, the AP gain of pedestrians is hence small or even nonexistent. From Table 1, RadarNet\* shows +1.9% AP for group (1) and +0.0% AP for group (2) averagely over the CenterPoint, while Bi-LRFusion shows +3.5% AP for group (1) and +1.2% AP for group (2) averagely. We summarize that RadarNet\* improves the AP values for group (1) objects, but not for group (2). The height missing problem in each Radar feature makes it difficult to detect specific tall objects. On the other hand, our Bi-LRFusion effectively avoids this problem and improves the AP for both group (1) and (2) over the baseline CenterPoint, regardless of the object’s height.

**Effect of Object Velocity.** Next, we look at the advantage of Bi-LRFusion over objects that move at different speeds. Following [4], we consider objects with speed  $> 0.5$  m/s as moving objects, 5m/s and 10m/s are the borderlines to distinguish low and medium moving state. More concretely, we also show the percentages of cars and motorcycles within different velocity ranges in the nuScenes dataset in Table 6: around 70% of cars and motorcycles are stationary, only around 4% of them are moving faster than

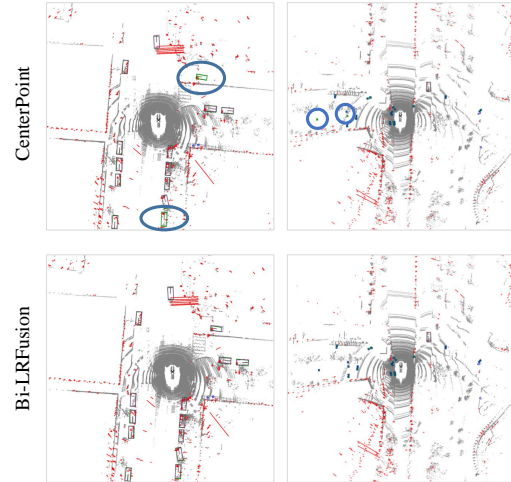


Figure 5. Qualitative comparison between the LiDAR-only method CenterPoint [40] and our Bi-LRFusion. The grey dots and red lines are LiDAR points and Radar points with velocities. We visualize the prediction and ground truth (green) boxes. Specifically, blue circles represent the missed detections from CenterPoint but Bi-LRFusion corrects them by integrating Radar data.

$> 10$  m/s, and the rest of them are uniformly distributed in the low / medium velocity range. Tables 5 (a) and (b) show the mAP and mAVE of the LiDAR-only detector and Bi-LRFusion for cars and motorcycles that move at different speeds. The results show that the performance improvement is more pronounced for objects with high velocity. This demonstrates the effectiveness of Radar in detecting dynamic objects, especially those with faster speeds.

## 5. Conclusion

In this paper, we introduce Bi-LRFusion, a bi-directional fusion framework that fuses the complementary LiDAR and Radar data for improving 3D dynamic object detection. Unlike existing LiDAR-Radar fusion schemes that directly integrate Radar features into the LiDAR-centric pipeline, we first make Radar features more discriminative by transferring the knowledge from LiDAR data (including both raw points and BEV features) to the Radar features, which alleviates the problems with Radar data, *i.e.*, lack of object height measurements and extreme sparsity. With our bi-directional fusion framework, Bi-LRFusion outperforms the earlier schemes for detecting dynamic objects on both nuScenes and ORR datasets.

**Acknowledgements.** We thank DequanWang for his help. This work was done during her internship at Shanghai AI Laboratory. This work was supported by the Chinese Academy of Sciences Frontier Science Key Research Project ZDBS-LY-JSC001, the Australian Medical Research Future Fund MRFAI000085, the National Key R&D Program of China(No.2022ZD0160100), and Shanghai Committee of Science and Technology(No.21DZ1100100).



## References

- [1] Dan Barnes, Matthew Gadd, Paul Murcutt, Paul Newman, and Ingmar Posner. The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6433–6438. IEEE, 2020. [2](#)
- [2] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11682–11692, 2020. [7](#)
- [3] Stefan Briskén, Florian Ruf, and Felix Höhne. Recent evolution of automotive imaging radar and its information content. *IET Radar, Sonar & Navigation*, 12(10):1078–1081, 2018. [2](#)
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. [2](#), [3](#), [6](#), [8](#)
- [5] Sarah H Cen and Paul Newman. Radar-only ego-motion estimation in difficult settings via graph matching. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 298–304. IEEE, 2019. [6](#)
- [6] Qi Chen, Lin Sun, Ernest Cheung, and Alan L Yuille. Every view counts: Cross-view consistency in 3d object detection with hybrid-cylindrical-spherical voxelization. *Advances in Neural Information Processing Systems*, 33:21224–21235, 2020. [7](#)
- [7] Yuwei Cheng, Hu Xu, and Yimin Liu. Robust small object detection on the water surface through fusion of camera and millimeter wave radar. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15263–15272, 2021. [3](#)
- [8] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1201–1209, 2021. [2](#), [5](#)
- [9] Jiajun Deng, Wengang Zhou, Yanyong Zhang, and Houqiang Li. From multi-view to hollow-3d: Hallucinated hollow-3d r-cnn for 3d object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(12):4722–4734, 2021. [1](#), [2](#)
- [10] Yifan Duan, Jie Peng, Yu Zhang, Jianmin Ji, and Yanyong Zhang. Pfilter: Building persistent maps through feature filtering for fast and accurate lidar-based slam, 2022. [6](#)
- [11] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8458–8468, 2022. [2](#)
- [12] Adam W Harley, Zhaoyuan Fang, Jie Li, Rares Ambrus, and Katerina Fragkiadaki. A simple baseline for bev perception without lidar. *arXiv preprint arXiv:2206.07959*, 2022. [3](#)
- [13] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. [3](#)
- [14] Youngseok Kim, Jun Won Choi, and Dongsuk Kum. Grifnet: Gated region of interest fusion network for robust 3d object detection from radar point cloud and monocular image. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10857–10864. IEEE, 2020. [3](#)
- [15] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. [2](#), [4](#), [5](#), [6](#), [7](#)
- [16] Xin Li, Botian Shi, Yuenan Hou, Xingjiao Wu, Tianlong Ma, Yikang Li, and Liang He. Homogeneous multi-modal feature fusion and interaction for 3d object detection. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVIII*, pages 691–707. Springer, 2022. [1](#)
- [17] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022. [3](#)
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [6](#)
- [19] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022. [5](#)
- [20] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3164–3173, 2021. [2](#)
- [21] Ramin Nabati and Hairong Qi. Centerfusion: Center-based radar and camera fusion for 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1527–1536, 2021. [1](#)
- [22] Felix Nobis, Ehsan Shafiei, Phillip Karle, Johannes Betz, and Markus Lienkamp. Radar voxel fusion for 3d object detection. *Applied Sciences*, 11(12):5598, 2021. [1](#), [3](#), [6](#)
- [23] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. [2](#), [4](#), [5](#)
- [24] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. [2](#), [4](#)
- [25] Kun Qian, Shilin Zhu, Xinyu Zhang, and Li Erran Li. 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*, pages 444–453, 2021. 1, 3, 6, 7
- [26] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pvrnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. *arXiv preprint arXiv:2102.00463*, 2021. 2
- [27] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 770–779, 2019. 2
- [28] Liat Sless, Bat El Shlomo, Gilad Cohen, and Shaul Oron. Road scene understanding by occupancy grid learning from sparse radar clusters using semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1
- [29] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4604–4612, 2020. 7
- [30] Dequan Wang, Yifan Duan, Xiaoran Fan, Chengzhen Meng, Jianmin Ji, and Yanyong Zhang. Maroam: Map-based radar slam through two-step feature selection. *arXiv preprint arXiv:2210.13797*, 2022. 6
- [31] Jun Wang, Shiyi Lan, Mingfei Gao, and Larry S Davis. Infocus: 3d object detection for autonomous driving with dynamic information modeling. In *European Conference on Computer Vision*, pages 405–420. Springer, 2020. 7
- [32] Yingjie Wang, Qiuyu Mao, Hanqi Zhu, Yu Zhang, Jianmin Ji, and Yanyong Zhang. Multi-modal 3d object detection in autonomous driving: a survey. *arXiv preprint arXiv:2106.12735*, 2021. 1
- [33] Yue Wang and Justin M Solomon. Object dgcnn: 3d object detection using dynamic graphs. *Advances in Neural Information Processing Systems*, 34:20745–20758, 2021. 2
- [34] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. 2
- [35] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 2, 4, 6
- [36] Bin Yang, Runsheng Guo, Ming Liang, Sergio Casas, and Raquel Urtasun. Radarnet: Exploiting radar for robust perception of dynamic objects. In *European Conference on Computer Vision*, pages 496–512. Springer, 2020. 1, 3, 5, 6, 7
- [37] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018. 7
- [38] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11040–11048, 2020. 2
- [39] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1951–1960, 2019. 2
- [40] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11784–11793, 2021. 1, 2, 5, 6, 7, 8
- [41] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. 2
- [42] Xinge Zhu, Yuexin Ma, Tai Wang, Yan Xu, Jianping Shi, and Dahua Lin. Ssn: Shape signature networks for multi-class object detection from point clouds. In *European Conference on Computer Vision*, pages 581–597. Springer, 2020. 7