

# RCTrans: Radar-Camera Transformer via Radar Densifier and Sequential Decoder for 3D Object Detection

Yiheng Li<sup>1,2,\*</sup>, Yang Yang<sup>1,2,\*</sup>, Zhen Lei<sup>1,2,3,†</sup>

<sup>1</sup> State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS),  
Institute of Automation, Chinese Academy of Sciences

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>3</sup> Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science & Innovation,  
Chinese Academy of Sciences

liyiheng2024@ia.ac.cn, yang.yang@nlpr.ia.ac.cn, zhen.lei@ia.ac.cn

## Abstract

In radar-camera 3D object detection, the radar point clouds are sparse and noisy, which causes difficulties in fusing camera and radar modalities. To solve this, we introduce a novel query-based detection method named Radar-Camera Transformer (RCTrans). Specifically, we first design a Radar Dense Encoder to enrich the sparse valid radar tokens, and then concatenate them with the image tokens. By doing this, we can fully explore the 3D information of each interest region and reduce the interference of empty tokens during the fusing stage. We then design a Pruning Sequential Decoder to predict 3D boxes based on the obtained tokens and random initialized queries. To alleviate the effect of elevation ambiguity in radar point clouds, we gradually locate the position of the object via a sequential fusion structure. It helps to get more precise and flexible correspondences between tokens and queries. A pruning training strategy is adopted in the decoder, which can save much time during inference and inhibit queries from losing their distinctiveness. Extensive experiments on the large-scale nuScenes dataset prove the superiority of our method, and we also achieve new state-of-the-art radar-camera 3D detection results.

**Code** — <https://github.com/liyih/RCTrans>

## 1 Introduction

3D object detection is an important perceptual task that can be widely applied in many fields. Existing high-precision detection algorithms often rely on the input of LiDAR (Yan, Mao, and Li 2018; Yin, Zhou, and Krahenbuhl 2021). However, LiDAR is expensive and easily damaged (Lin et al. 2024), which is not conducive to the commercialization of the algorithm. To reduce the practical cost, recent works begin to focus on using low-cost radars and multi-view cameras to improve detection performance. Camera sensors can provide rich texture information, while radar sensors can provide three-dimensional information that is not affected by different weather conditions (Kim et al. 2023a).

\*These authors contributed equally.

†Corresponding author.

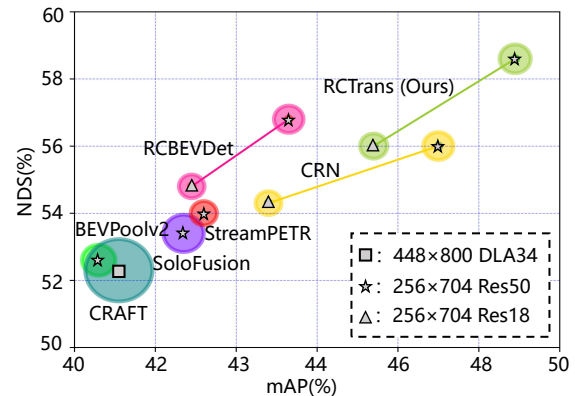


Figure 1: Comparison between RCTrans and other camera or radar-camera 3D detection methods. The size of the circle represents the time consumption of inference. All experiments are conducted on nuScenes val set and the speed is measured on a single NVIDIA RTX3090 GPU.

However, the radar sensor has two inherent drawbacks, i.e., sparsity and noise. Firstly, a radar typically releases mmWave into the physical space, which hits the object and then is reflected to the receiver. Due to the lack of diffused reflection, only a small part of objects can be reflected to the radar receiver (Singh et al. 2023). The number of non-empty radar pillars is only approximately 10% of average numbers in LiDAR pillars (Bang et al. 2024). Secondly, the radar used in autonomous driving (such as in nuScenes (Caesar et al. 2020)) cannot obtain an accurate height of objects, which can bring the noise in azimuth (x) and depth (z) components (Singh et al. 2023). In addition, some noise may also be led to the axes when the beam-width of the mmWave radar is too large (Singh et al. 2023). The sparsity and noise of radar data will bring challenges to the cross-modality fusion. Specifically, sparsity causes too much invalid radar representation, while noise makes wrong or inaccurate alignments between camera and radar representation.

As shown in Fig. 1, we compare RCTrans with existing models on nuScenes val set. Previous methods may overlook the sparsity of the radar (Kim et al. 2023b,a), or need the radar cross-section (RCS) as the prior (Lin et al.

2024). They cannot adaptively fill every empty BEV grid. Moreover, the SOTA radar-camera solutions (Kim et al. 2023b; Lin et al. 2024) adopt bird’s eye view (BEV) (Huang et al. 2021) as the unified space to fuse features. Such kind of methods require strict correspondences between multi-modality features in BEV space, but it is difficult to establish precise correspondences for radar with positional offset. Inspired by the success of query-based 3D detection (Liu et al. 2022; Yan et al. 2023), we use the query-based paradigm to establish interaction between object queries and multi-modality tokens. A positional embedding is used to construct implicit associations, allowing the model to adaptively learn more flexible correspondences.

In this paper, we propose a novel Radar-Camera Transformer (RCTrans) framework via radar densifier and sequential decoder to overcome the sparsity and noise of radar inputs. To be specific, we design a Radar Dense Encoder (RDE) to avoid blurring some detailed information while densifying every empty BEV grid. It uses a downsample-then-upsample structure and obtains multi-scale information through a skip connection. Self-attention is adopted at the smallest BEV scale to adaptively fill each grid. Considering that information from different modalities can establish more flexible interaction with object queries, we design the Pruning Sequential Decoder, which gradually fuses multi-modality features and locates the position of each object in a step-by-step way. In each layer, two separate transformers will be used to fuse the radar and image information, respectively. This approach allows the model to establish the connections between different modalities and object queries based on their distinct traits. Thereafter, we predict the new position of queries and recompute position embedding after each decoder layer. In the next layer, we can use the updated position embedding to calculate more precise correspondences and obtain more accurate fusion results. A pruning training strategy is used to speed up the inference and prevent object queries from losing their distinctiveness.

We conduct extensive experiments on the large-scale autonomous driving dataset nuScenes (Caesar et al. 2020), which demonstrates the effectiveness of our proposed RCTrans. Our contributions are summarized as:

- We introduce a new query-based 3D detection framework named Radar-Camera Transformer (RCTrans), which can solve the problems caused by the sparsity and noise of radar point clouds.
- We first design the Radar Dense Encoder to densify the sparse radar features, and then propose the Pruning Sequential Decoder to make flexible alignment and effective fusion in the prediction process.
- We confirm the efficacy of RCTrans by achieving superior 3D detection performance in the nuScenes dataset with an acceptable speed. What’s more, the ablation study proves the role of each proposed module.

## 2 Related Work

### 2.1 BEV-Based 3D Object Detection

Bird’s-eye-view (BEV) is a dense 2D representation, where each grid corresponds to a certain region in the scene (Mao

et al. 2023). LSS (Phillion and Fidler 2020) innovatively solves the transformation problem from multi-view cameras to BEV features. BEVDet (Huang et al. 2021) is a pioneer in applying BEV to 3D object detection. The major difficulty of building BEV lies in inaccurate depth estimation, and BEVDepth (Li et al. 2023b) solves this bottleneck by using additional depth supervision. In order to predict more accurate speed and reduce the missed detection of objects, more and more works try to incorporate temporal information. BEVFormer (Li et al. 2022) is the first to introduce sequential temporal modeling in multi-view 3D object detection. SOLOFusion (Park et al. 2022) proposes a stream video paradigm to make long-term temporal modeling. BEV can also serve as a unified representation space for the fusion of multi-modality features. BEVFusion (Liu et al. 2023c; Liang et al. 2022) is a pioneer in choosing BEV as a unified representation space for fusion. BEV can provide explicit correspondence for different modalities, which is extremely convenient for fusion. However, alignment errors between different modal BEVs can affect the fusion effect. Some approaches (Song et al. 2024; Lin et al. 2024) try to solve this issue via a deformable module.

### 2.2 Query-Based 3D Object Detection

The query-based method is another important paradigm in 3D object detection. In these methods, object queries typically interact with the tokens of input data. Inspired by the successful query-based methods in 2D detection (Carion et al. 2020), DETR3D (Wang et al. 2022a) introduces a set of points to be the references, each of them represents an object query. Later, researchers begin to focus on how to effectively extract features from multi-view images to the object queries. PETR (Liu et al. 2022) proposes 3D position-aware position embedding for image features in the LiDAR coordinate. CAPE (Xiong et al. 2023) constructs the position embedding in camera coordinates to reduce the influence of variant camera extrinsic. StreamPETR (Wang et al. 2023) proposes a query-based long-temporal fusion algorithm that can transfer historical information from objects to the current frame through object queries. PETRv2 (Liu et al. 2023b) aggregates short-temporal information by transforming the previous 3D coordinates to the current coordinate system. Some methods also explore the query-based multi-modality fusion paradigm. Futr3d (Chen et al. 2023) introduces the first unified end-to-end sensor fusion framework in 3D detection. CMT (Yan et al. 2023) proposes a 3D position embedding to fuse multi-modality tokens.

### 2.3 Radar-Camera 3D Object Detection

To explore low-cost 3D perception solutions, more and more methods have begun to focus on using radar sensors. RadarNet (Yang et al. 2020) utilizes radar information to complement other sensors in the form of Doppler velocity. Based on the frustum-based association method, CenterFusion (Nabati and Qi 2021) accurately associates radar detections to objects in the image. What’s more, it adopts a middle-fusion approach to complement the image features via radar feature maps. CRAFT (Kim et al. 2023a) introduces the Spatio-Contextual Fusion Transformer to conduct

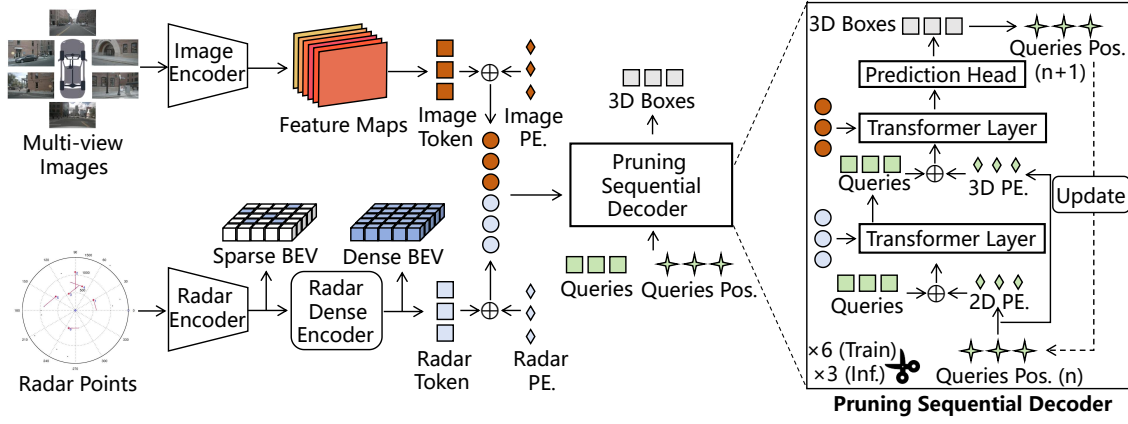


Figure 2: The overall architecture of Radar-Camera Transformer. RCTrans can be divided into two key components: (1) Token generator which extracts multi-modality tokens and sums them with corresponding position embedding and (2) Pruning Sequential Decoder which updates the randomly initialized queries and queries position in each layer and makes more precise fusion in a step-by-step way. The final predicted results are the outputs of the last layer after pruning.

radar-camera detection. CRN (Kim et al. 2023b) transforms the image feature into BEV space with the help of radar points. RADIANT (Long et al. 2023) predicts the object center offsets and associates these offsets with the image detector. It then obtains accurate depth estimates for object detection and enhances camera features with estimated depth. RCFusion (Zheng et al. 2023) proposes a weighted fusion module to fuse radar and camera BEV features. RCBEVDet (Lin et al. 2024) fuses and predicts in the BEV space and proposes the RadarBEVNet to better extract radar features.

In contrast, we introduce a query-based fusion method named RCTrans, which focuses on how to solve the inherent drawbacks of the radar sensor and achieves high results.

### 3 Method

The pipeline of our proposed Radar-Camera Transformer (RCTrans) is shown in Fig. 2. Firstly, two parallel branches are used to extract multi-modality tokens. In the radar branch, we will use the Radar Dense Encoder to densify the valid radar features. Then, position embeddings of multi-modality tokens are generated and added to the tokens. Finally, the randomly initialized queries will be sent to the Pruning Sequential Decoder along with the tokens to predict 3D boxes. The entire training is an end-to-end process and does not require freezing any parameters.

#### 3.1 Token Generator

For the camera branch, given the images  $I \in R^{N_i \times H \times W \times 3}$  from  $N_i$  views where  $H$  and  $W$  are the size of the image, the image encoder is used to encode the images into feature maps  $F \in R^{N_i \times \frac{H}{16} \times \frac{W}{16} \times C}$ . Then, the  $F$  is flattened into the image tokens. For the radar branch, we follow Futr3d (Chen et al. 2023) to use the 3D coordinate, compensated velocities, and the time offset as the input channel of the radar points  $P \in R^{N_r \times 5}$ , where  $N_r$  is the number of the points. After that, we use MLP to obtain the features of each pillar and scatter them to the corresponding BEV grids. Due to

the sparsity of radar points, more than 90% BEV grids are empty, which may hurt the following detection results. To solve it, we introduce a Radar Dense Encoder to adaptively fill every grid while also keeping multi-scale information.

**Radar Dense Encoder.** One naive solution is to directly increase the number of the BEV encoder layers. However, this may smooth the small objects in the background. Moreover, commonly used BEV encoders, such as SECOND (Yan, Mao, and Li 2018), are mainly designed for LiDAR and not entirely suitable for radar inputs. RCBEVDet (Lin et al. 2024) proposes to use the radar cross-section (RCS) to determine the scattering radius of radar points. However, the scattered results rely on the quality of RCS and the number of scattered features in each BEV grid is different. We aim to find a network with a simple structure, that can adaptively fill each BEV grid and aggregate multi-scale information to facilitate detection of objects with different sizes.

The simplest method to adaptively fill each BEV grid is to use the global self-attention mechanism. However, due to the large number of BEV grids, this method consumes a long time. To solve it, we first downsample the original BEV features and interact with them in a minimal resolution. At the same time, the downsampling process can greatly reduce the number of invalid grids, thereby improving the quality of features obtained after interaction. Inspired by U-net (Ronneberger, Fischer, and Brox 2015), we connect features of different receptive field sizes at the same BEV resolution, which effectively preserves the features of objects of different sizes. The architecture of the Radar Dense Encoder (RDE) is shown in Fig. 3. Specifically, given the radar BEV feature  $B \in R^{H_r \times W_r \times C_r}$ , where  $H_r$  and  $W_r$  are the size of the BEV map, we first extract multi-level downsampled features  $B_d = \{B_i \in R^{\frac{H_r}{2^i} \times \frac{W_r}{2^i} \times C_i}, i = 1, 2, 3\}$ . Then, we conduct self-attention in  $B_3$  and obtain adaptively filled BEV map  $B_f$ . During this process, each grid is added by the 2D position embedding. After that, we upsample the  $B_f$  to the same size as the original  $B$ . Skip connection is used to




Figure 3: The overall architecture of Radar Dense Encoder. RDE uses a downsample-then-upsample architecture. Self-attention is used to fill the empty grids in the minimal resolution. Skip connection is used in the upsampling process.

fuse the upsampled BEV map and the same size features of  $B_d$  in the upsampling process.

**Position Embedding.** For image position embedding, we use the 3D position embedding proposed in PETR (Liu et al. 2022). Given a image token  $T_i$ , a series points  $p(u, v) = \{p_i(u, v) = (u \times d_i, v \times d_i, d_i, 1), i = 1, 2, \dots, d\}$  is pre-defined in the camera frustum space. Here,  $u$  and  $v$  are the indices of the token in image space, and  $d$  is the number of points along the depth axis. After that, the image position embedding is calculated via Eq. 1.

$$PE_{im} = \Phi_{im}(Kp(u, v)), \quad (1)$$

where  $K$  is the transformation matrix that transforms camera frustum to 3D world space.  $\Phi_{im}(\cdot)$  is the MLP function. As the radar can not obtain precise height information of objects, for radar position embedding, we use the 2D BEV embedding, ignoring the height information of the BEV grid. The radar position embedding is calculated via Eq. 2.

$$PE_{ra} = \Phi_{ra}(\Psi(h, w)), \quad (2)$$

where  $(h, w)$  is the 2D coordinate of BEV grid,  $\Phi_{ra}(\cdot)$  is the MLP function, and  $\Psi(\cdot)$  is the sine-cosine function. Through positional embedding, information from different modalities can be implicitly aligned with object queries in 3D space.

### 3.2 Pruning Sequential Decoder

The Pruning Sequential Decoder consists of multiple layers, in each layer, we fuse the image tokens with object queries through  $PE_{im}$  and  $PE_{3d}$ , and fuse the radar tokens with object queries through  $PE_{ra}$  and  $PE_{2d}$  due to the inaccurate height information of radars. Following previous query-based methods (Liu et al. 2022; Wang et al. 2022b), we first initialize  $n$  learnable references  $R = \{r_i = (r_{xi}, r_{yi}, r_{zi}), i = 1, \dots, n\}$  in 3D scene to be the position of the object queries, which uniformly distribute between  $[0, 1]$ . The query feature  $F_q$  corresponding to each reference is initialized as a vector of all zeros.

**Query Feature Update.** The query positions  $R$  are first projected into radar and image space. The projected query positions  $R_{ra}$  in radar space can be obtained by Eq. 3.

$$\begin{cases} r'_{xi} = r_{xi} \times (x_{max} - x_{min}) + x_{min} \\ r'_{yi} = r_{yi} \times (y_{max} - y_{min}) + y_{min} \\ r'_{zi} = r_{zi} \times (z_{max} - z_{min}) + z_{min}, \end{cases} \quad (3)$$

where  $\{[x_{max}, x_{min}], [y_{max}, y_{min}], [z_{max}, z_{min}]\}$  is the

range of valid 3D world space. The projected query positions  $R_{im}$  in image space can be obtained by Eq. 4.

$$R_{im} = K^{-1}R_{ra}, \quad (4)$$

where  $K$  establishes the transformation from camera frustum to 3D world space. Following CMT (Yan et al. 2023), we adopt the shared encoder used in image/radar position embedding generation to encode 3D/2D position embedding of queries. The  $PE_{3d}$  and  $PE_{2d}$  are obtained by Eq. 5.

$$PE_{3d} = \Phi_{im}(R_{im}), PE_{2d} = \Phi_{ra}(\Psi(R_{ra})) \quad (5)$$

Different from the common decoder used for multi-modality tokens which concatenates them and uses a large layer to fuse them together, we propose to split this large layer into two small ones and fuse the multi-modality tokens separately. We name it sequential structure. The updated queries are calculated by Eq. 6.

$$F_q^{n+1} = f_2(f_1(F_q^n + PE_{2d}, T_r) + PE_{3d}, T_i), \quad (6)$$

where  $T_r$  and  $T_i$  represent radar and image tokens.  $f_1(\cdot, \cdot)$  and  $f_2(\cdot, \cdot)$  represent the transformer layer.

**Query Position Update.** At the end of each decoder layer, we will predict the positions of the queries. In the next layer, new position embedding will be generated based on the updated positions. Given the updated query  $F_q^{n+1}$ , we predict the offset  $\Delta R$  of the query positions, and the updated position  $R_{n+1}$  can be calculated by  $R_n + \Delta R$ .

**Pruning Training Strategy.** In each layer, we use a sequential structure to fuse multi-modality information, which results in using twice as many transformer layers as traditional decoders. This can cause additional inference time. What's more, as we update the positions of queries after each decoder layer, some object queries may gradually be located in the same region and lose the feature distinctiveness. This causes information from some regions to be overlooked, and the attention mechanism may fail to learn effective representation learning concepts which prevents the model from achieving the expected performance improvement (Zhou et al. 2021). To this end, we propose a pruning training strategy, which uses 6 layers of decoders during training and only 3 layers of decoders during inference.

### 3.3 Training Head and Task Loss

For the 3D object prediction head, we follow PETR (Liu et al. 2022) to use the FFN to regress each attribute of the object and conduct deep supervision. For the 3D object

Method	Input	Image Backbone	Image Size	NDS (%) $\uparrow$	mAP (%) $\uparrow$	Latency (ms) $\downarrow$
CRN (2023b)	RC	ResNet18	$256 \times 704$	54.3	44.8	35.8
RCBEVDet (2024)	RC	ResNet18	$256 \times 704$	54.8	42.9	<b>35.3</b>
RCTrans(ours)	RC	ResNet18	$256 \times 704$	56.0	47.4	39.8
BEVDepth (2023b)	C	ResNet50	$256 \times 704$	47.5	35.1	86.2
BEVPoolv2 (2022b)	C	ResNet50	$256 \times 704$	52.6	40.6	60.2
SOLOFusion (2022)	C	ResNet50	$256 \times 704$	53.4	42.7	87.7
StreamPETR (2023)	C	ResNet50	$256 \times 704$	54.0	43.2	36.9
CRN (2023b)	RC	ResNet50	$256 \times 704$	56.0	49.0	49.0
RCBEVDet (2024)	RC	ResNet50	$256 \times 704$	56.8	45.3	47.0
RCTrans(ours)	RC	ResNet50	$256 \times 704$	58.6	50.9	52.2
RCBEV4d (2022a)	RC	Swin-T	$256 \times 704$	49.7	38.1	-
RCBEVDet (2024)	RC	Swin-T	$256 \times 704$	56.2	49.6	54.9
RCTrans(ours)	RC	Swin-T	$256 \times 704$	<b>59.4</b>	<b>52.0</b>	60.5

Table 1: 3D object detection performance comparison on nuScenes val set with different image backbones. ‘C’ and ‘R’ represent camera and radar, respectively.

tracking, we follow CenterPoint (Yin, Zhou, and Krahenbuhl 2021) to use the velocity-based greedy matching algorithm. This is an offline reasoning process and does not require additional training. For the supervision of the model, we adopt focal loss (Lin et al. 2017) for classification and L1 loss for 3D bounding box regression. We utilize the Hungarian algorithm to assign the predicted objects and the ground-truths. The loss can be calculated by Eq. 7.

$$L(o, \hat{o}) = \omega_1 L_{cls}(c, \hat{c}) + \omega_2 L_{reg}(b, \hat{b}), \quad (7)$$

where  $\omega_1$  and  $\omega_2$  are the hyper-parameters to balance classification and regression loss. The variables  $c$  and  $b$  are the predicted classification and bounding box, while  $\hat{c}$  and  $\hat{b}$  are the corresponding ground truths. This kind of supervision is used for the predicted 3D boxes after each decoder layer.

## 4 Experiments

### 4.1 Datasets and Metrics

Following previous approaches (Kim et al. 2023a,b; Lin et al. 2024), we conduct comprehensive experiments on a large-scale autonomous driving dataset for 3D radar-camera object detection, nuScenes (Caesar et al. 2020). This dataset has a total of 1000 scenes and is officially divided into 700/150/150 for training/validation/testing. For each frame, nuScenes has 6 images and 5 radar point clouds covering  $360^\circ$ . There are around 1.4 million annotated 3D bounding boxes for ten classes. We select nuScenes detection score (NDS) and mean average precision (mAP) as metrics to evaluate the 3D detection. For NDS, it is the weighted sum of mAP and other official predefined metrics, including Average Translation Error (ATE), Average Scale Error (ASE), Average Orientation Error (AOE), Average Velocity Error (AVE), Average Attribute Error (AAE). For 3D tracking, we use the official multi-object tracking accuracy (AMOTA), average multi-object tracking precision (AMOTP), false positive (FP), false negative (FN), and ID switches (IDS) as the metrics.

### 4.2 Implementation Details

We implement RCTrans based on StreamPETR (Wang et al. 2023) and MMDetection3D (Contributors 2020) codebases. Following CRN (Kim et al. 2023b), we accumulate the information of 4 previous frames to the current frames. We use the object-centric temporal modeling proposed in StreamPETR to conduct temporal fusion. The number of decoder layers is set to 6 during training and 3 during inference. The output of the last layer after pruning is inserted into the memory queue in temporal fusion. We set the number of queries, memory queue, and propagated queries to 900, 512, and 128, respectively. For radar, we accumulate 6 previous radar sweeps following CRAFT (Kim et al. 2023a) and set the maximum number of radar points to 2048. The size of radar BEV is set to  $128 \times 128$ . We train our network for 90 epochs with a batch size of 32 on 8 NVIDIA A100 GPUs. The speed is evaluated on a single NVIDIA RTX3090 GPU. We conduct optimization based on AdamW with weight decay value  $10^{-2}$ , and the learning rate is adjusted by cycle policy with the initial value  $4 \times 10^{-4}$ .

### 4.3 Performance Comparison

**3D Object Detection.** We first compare our method with the state-of-the-art on nuScenes val set under different images backbones, including ResNet18 (He et al. 2016), ResNet50, and Swin-T (Liu et al. 2021). As shown in Table 1, our method achieves the best performance under different backbones with a slight increase in time consumption. For example, when using Swin-T as the backbone and setting the image size to  $256 \times 704$ , compared to the SOTA radar-camera solution RCBEVDet, RCTrans improves NDS by 3.2% and mAP by 2.4%, while the latency increases around 5 ms. The great experimental results of multiple backbones demonstrate that RCTrans has excellent and strong adaptability, which is very beneficial for model deployment and migration in practical applications. What’s more, RCTrans beats all the camera-only methods, including our camera stream baseline StreamPETR, which proves our method can effectively use radar information to supplement the detection re-

Method	Input	Backbone	mATE	mASE	mAOE	mAVE	mAAE	NDS $\uparrow$	mAP $\uparrow$
Radar-PointGNN (2021)	R	-	1.024	0.859	0.897	1.020	0.931	3.4	0.5
KPConvPillars (2022)	R	-	0.823	0.428	0.607	2.081	1.000	13.9	4.9
BEVFormer(2022)	C	V2-99	0.582	0.256	0.375	0.378	0.126	56.9	48.1
PETR-v2 (2023b)	C	V2-99	0.561	0.243	0.361	0.343	0.120	58.2	49.0
BEVDepth (2023b)	C	ConvNeXt-B	0.445	0.243	0.352	0.347	0.127	60.9	52.0
BEVStereo (2023a)	C	ConvNeXt-B	0.431	0.246	0.358	0.357	0.138	61.0	52.5
SOLOFusion (2022)	C	ConvNeXt-B	0.453	0.257	0.376	0.276	0.148	61.9	54.0
SparseBEV (2023a)	C	V2-99	0.485	0.244	0.332	0.246	0.117	63.6	55.6
StreamPETR (2023)	C	V2-99	0.493	0.241	0.343	0.243	0.123	63.6	55.0
CenterFusion (2021)	RC	DLA34	0.631	0.261	0.516	0.615	0.115	44.9	32.6
RCBEV (2023)	RC	Swin-T	0.484	0.257	0.587	0.702	0.140	48.6	40.6
MVFusion (2023)	RC	V2-99	0.569	0.246	0.379	0.781	0.128	51.7	45.3
CRAFT (2023a)	RC	DLA34	0.467	0.268	0.456	0.519	0.114	52.3	41.1
CRN (2023b)	RC	ConvNeXt-B	0.416	0.264	0.456	0.365	0.130	62.4	57.5
RCBEVDet (2024)	RC	V2-99	0.390	0.234	0.362	0.259	0.113	63.9	55.0
RCTrans(ours)	RC	V2-99	0.459	0.245	0.392	0.198	0.121	<b>64.7</b>	<b>57.8</b>

Table 2: 3D object detection performance comparison on nuScenes test set. The V2-99 (Lee and Park 2020) backbone is pre-trained on external dataset DDAD (Guizilini et al. 2020).

Method	Input	AMOTA $\uparrow$	AMOTP $\downarrow$	FP $\downarrow$	FN $\downarrow$	IDS $\downarrow$
DEFT(2021)	C	17.7	1.564	22163	60565	6901
QD-3DT(2022)	C	21.7	1.550	16495	60156	6856
CC-3DT(2022)	C	41.0	1.274	18114	42910	3334
Sparse4D(2022)	C	51.9	1.078	19626	32954	1090
CRN(2023b)	RC	56.9	<b>0.809</b>	16822	41093	946
RCTrans	RC	<b>59.6</b>	0.969	<b>14700</b>	<b>29491</b>	<b>654</b>

Table 3: 3D object tracking performance comparison on nuScenes test set. We conduct experiments based on the V2-99 backbone.

sults. We also compare our method with the state-of-the-art radar, camera, and radar-camera approaches on nuScenes test set. As shown in Table 2, our proposed RCTrans outperforms all the competitive methods by achieving 64.7% NDS and 57.8% mAP. Notably, RCTrans receives a giant performance improvement in mAVE, which proves that our model extracts useful velocity compensation information from radar data. Qualitative results are shown in Fig. 4.

**3D Object Tracking.** We further prove the generalization ability of RCTrans on 3D object tracking task. As shown in Table 3, we compare RCTrans with the existing camera and radar-camera tracking solutions on nuScenes test set. Overall, our approach yields the best results. Compared to CRN, our method significantly improves AMOTA, FP, FN, and IDS. The tracking results we compared are all obtained based on velocity-based closest distance matching in CenterPoint (Yin, Zhou, and Krahenbuhl 2021), so the improvement in tracking performance is mainly due to the more accurate prediction of the speed by our approach.

#### 4.4 Ablation Study

We conduct the ablation study on nuScenes val set to demonstrate the effectiveness of each component. All of our experiments are conducted under the condition of using ResNet50 as the backbone and setting the image size

Component	Details	NDS	mAP
Modality	radar $\Delta$	12.3	1.6
	camera $\nabla$	49.3	38.9
	radar+camera	51.7( $\uparrow$ 2.4)	43.0( $\uparrow$ 4.1)
	radar+camera $\diamond$	55.1( $\uparrow$ 3.4)	46.6( $\uparrow$ 3.6)
RDE	+U-structure	54.2( $\uparrow$ 2.5)	44.6( $\uparrow$ 1.6)
	+skip connection	56.3( $\uparrow$ 2.1)	47.4( $\uparrow$ 2.8)
	+self-attention	56.4( $\uparrow$ 0.1)	48.2( $\uparrow$ 0.8)
PSD	+position update	57.2( $\uparrow$ 0.8)	49.7( $\uparrow$ 1.5)
	+sequential structure	58.4( $\uparrow$ 1.2)	50.3( $\uparrow$ 0.6)
	+pruning strategy	58.6( $\uparrow$ 0.2)	50.9( $\uparrow$ 0.6)

Table 4: Ablation study on nuScenes val set.  $\Delta$ : the normal radar stream does not use any encoder on the BEV feature.  $\nabla$ : the camera stream is modified based on StreamPETR. We simplify it by removing the spatial alignment module and change the query position embedding.  $\diamond$ : using SECOND as the BEV encoder.

Training layer	Inference layer	NDS	mAP	latency(ms)
6	6	58.4	50.3	69.9
6	5	58.5	50.5	64.0
6	4	<b>58.7</b>	<b>51.0</b>	58.1
6	3	58.6	50.9	<b>52.2</b>
3	3	56.8	49.2	52.2

Table 5: Ablation study of pruning training strategy.

to  $256 \times 704$ . As shown in Table 4, each component can consistently improve performance. Compared to the single-modality stream, using multi-modality inputs can significantly achieve improvement. Compared to commonly used BEV encoders, such as SECOND (Yan, Mao, and Li 2018), using our proposed Radar Dense Encoder (RDE) can obtain 1.3% NDS and 1.6% mAP improvement. What’s more, the Pruning Sequential Decoder (PSD) increases NDS by 2.2% and mAP by 2.7%. We also conduct experiments on the

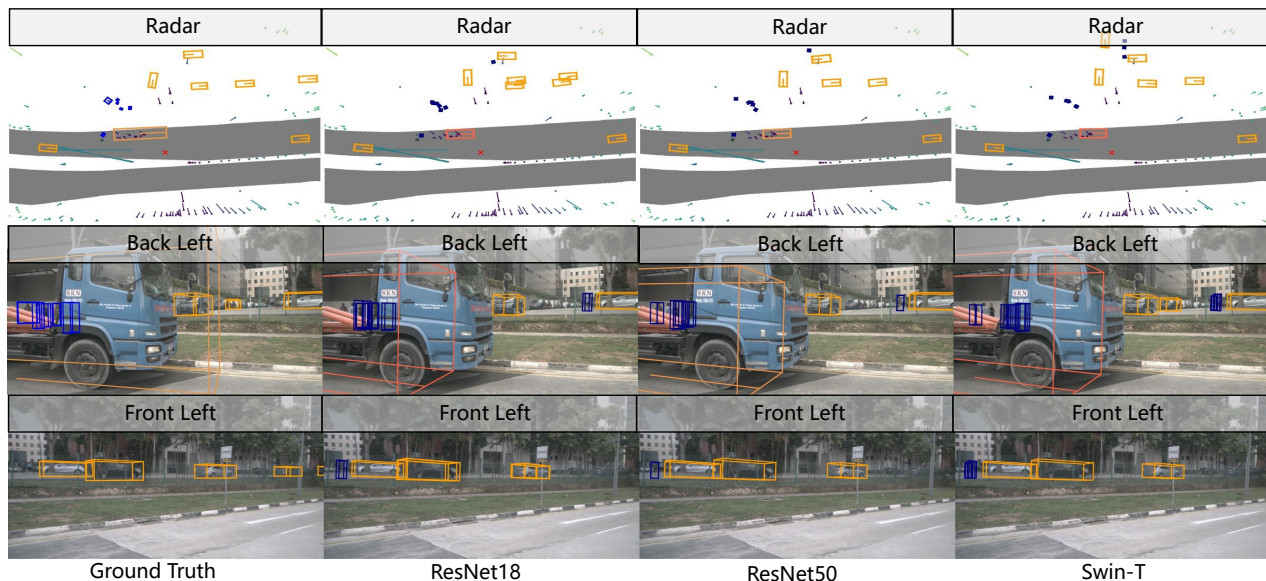


Figure 4: Qualitative results of RCTrans on nuScenes val set with different backbones. Due to space constraints, we only show the corresponding front left and back left images where objects are mainly distributed in this scene.

	Input	Drop	# of view drops			
			0	1	3	6
BEVDepth	C	C	49.4	41.1	24.2	0
CenterPoint	R	R	30.6	25.3	14.9	0
BEVFusion	C+R	C	63.9	58.5	45.7	14.3
		R	59.9	50.9	34.4	
CRN	C+R	C	68.8	62.4	48.9	12.8
		R	64.3	57.0	43.8	
RCBEVDet	C+R	C	72.5	66.9	53.5	16.5
		R	71.6	66.1	62.1	
RCTrans	C+R	C	<b>73.4</b>	<b>71.9</b>	<b>68.8</b>	<b>20.3</b>
		R	<b>72.2</b>	<b>71.6</b>	<b>62.5</b>	

Table 6: Robustness analysis via car class mAP. When drops=6, it means that the inputs contain single modality.

pruning training strategy as shown in Table 5. We find that setting the number of decoders to 6 during training and only using the first 3 layers during inference results in great performance, reducing inference time by 17.7ms while increasing NDS by 0.2%. The reason for this phenomenon may be that recalculating the position embedding after each layer can quickly obtain accurate alignment and faster converge but make some queries focus on the same region. Moreover, fusing different modal information in each layer by sequential structure will result in a large total number of transformers in the decoder. Reducing the number of decoder layers during inference can speed up the inference procedure while the performance is not reduced.

#### 4.5 Robustness Analysis

To simulate the sensor malfunction in real autonomous driving scenarios, we randomly drop images or radar inputs and evaluate the results. For a fair comparison, following CRN

(Kim et al. 2023b) and RCBEVDet (Lin et al. 2024), we re-train the methods with data drop augmentation proposed in Autoalignv2 (Chen et al. 2022), which randomly drops 3 image inputs during training. During the test process, the number of the dropped sensors is set to 1, 3, and 6. The removed inputs will be replaced by a fully zero tensor of the same shape. Car class mAP is used to evaluate the robustness of the methods. As shown in Table 6, RCTrans outperforms all the existing methods under different sensor failure cases. Specifically, compared to RCBEVDet, RCTrans improves mAP by 5.0%, 15.3%, and 3.8% under these three camera failure conditions. Notably, when the number of dropped inputs is 3, RCTrans can still get the results close to the condition with null or 1 dropped input, which is an important improvement over past methods. The above phenomenon shows that our model can achieve a more stable effect under the sensor failure cases.

## 5 Conclusion

In this paper, we propose a Radar-Camera Transformer (RCTrans), a query-based method that can effectively solve the fusion difficulties caused by the sparsity and noise of radar inputs. To solve the problem of sparsity, we introduce the Radar Dense encoder which uses a downsample-then-upsample architecture. It can adaptively fill the invalid BEV grids while maintaining the information of small objects. To solve the problem of noise, we introduce the Pruning Sequential Decoder which uses a step-by-step strategy to gradually regress the objects. We also introduce the pruning training strategy for the decoder, which saves much time during inference while also preventing attention collapse. Extensive experiments demonstrate the effectiveness and robustness of RCTrans. We believe RCTrans can be an effective baseline to inspire future research.

## Acknowledgments

This work was supported in part by the Chinese National Natural Science Foundation Project 62206276, 62276254, U23B2054, and the InnoHK program.

## References

- Bang, G.; Choi, K.; Kim, J.; Kum, D.; and Choi, J. W. 2024. RadarDistill: Boosting Radar-based Object Detection Performance via Knowledge Distillation from LiDAR Features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15491–15500.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Chaabane, M.; Zhang, P.; Beveridge, J. R.; and O’Hara, S. 2021. Deft: Detection embeddings for tracking. *arXiv preprint arXiv:2102.02267*.
- Chen, X.; Zhang, T.; Wang, Y.; Wang, Y.; and Zhao, H. 2023. Futr3d: A unified sensor fusion framework for 3d detection. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 172–181.
- Chen, Z.; Li, Z.; Zhang, S.; Fang, L.; Jiang, Q.; and Zhao, F. 2022. Deformable feature aggregation for dynamic multi-modal 3D object detection. In *European conference on computer vision*, 628–644. Springer.
- Contributors, M. 2020. MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>.
- Fischer, T.; Yang, Y.-H.; Kumar, S.; Sun, M.; and Yu, F. 2022. Cc-3dt: Panoramic 3d object tracking via cross-camera fusion. *arXiv preprint arXiv:2212.01247*.
- Guizilini, V.; Ambrus, R.; Pillai, S.; Raventos, A.; and Gaidon, A. 2020. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2485–2494.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, H.-N.; Yang, Y.-H.; Fischer, T.; Darrell, T.; Yu, F.; and Sun, M. 2022. Monocular quasi-dense 3d object tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 1992–2008.
- Huang, J.; and Huang, G. 2022a. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*.
- Huang, J.; and Huang, G. 2022b. Bevpoolv2: A cutting-edge implementation of bevdet toward deployment. *arXiv preprint arXiv:2211.17111*.
- Huang, J.; Huang, G.; Zhu, Z.; Ye, Y.; and Du, D. 2021. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*.
- Kim, Y.; Kim, S.; Choi, J. W.; and Kum, D. 2023a. Craft: Camera-radar 3d object detection with spatio-contextual fusion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1160–1168.
- Kim, Y.; Shin, J.; Kim, S.; Lee, I.-J.; Choi, J. W.; and Kum, D. 2023b. Crn: Camera radar net for accurate, robust, efficient 3d perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17615–17626.
- Lee, Y.; and Park, J. 2020. Centermask: Real-time anchor-free instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13906–13915.
- Li, Y.; Bao, H.; Ge, Z.; Yang, J.; Sun, J.; and Li, Z. 2023a. Bevstereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1486–1494.
- Li, Y.; Ge, Z.; Yu, G.; Yang, J.; Wang, Z.; Shi, Y.; Sun, J.; and Li, Z. 2023b. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1477–1485.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 1–18. Springer.
- Liang, T.; Xie, H.; Yu, K.; Xia, Z.; Lin, Z.; Wang, Y.; Tang, T.; Wang, B.; and Tang, Z. 2022. Bevfusion: A simple and robust lidar-camera fusion framework. *Advances in Neural Information Processing Systems*, 35: 10421–10434.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, X.; Lin, T.; Pei, Z.; Huang, L.; and Su, Z. 2022. Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion. *arXiv preprint arXiv:2211.10581*.
- Lin, Z.; Liu, Z.; Xia, Z.; Wang, X.; Wang, Y.; Qi, S.; Dong, Y.; Dong, N.; Zhang, L.; and Zhu, C. 2024. RCBEVDet: Radar-camera Fusion in Bird’s Eye View for 3D Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14928–14937.
- Liu, H.; Teng, Y.; Lu, T.; Wang, H.; and Wang, L. 2023a. Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18580–18590.
- Liu, Y.; Wang, T.; Zhang, X.; and Sun, J. 2022. Petr: Position embedding transformation for multi-view 3d object detection. In *European Conference on Computer Vision*, 531–548. Springer.

- Liu, Y.; Yan, J.; Jia, F.; Li, S.; Gao, A.; Wang, T.; and Zhang, X. 2023b. Petrv2: A unified framework for 3d perception from multi-camera images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3262–3272.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu, Z.; Tang, H.; Amini, A.; Yang, X.; Mao, H.; Rus, D. L.; and Han, S. 2023c. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *2023 IEEE international conference on robotics and automation (ICRA)*, 2774–2781. IEEE.
- Long, Y.; Kumar, A.; Morris, D.; Liu, X.; Castro, M.; and Chakravarty, P. 2023. RADIANT: Radar-image association network for 3D object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1808–1816.
- Mao, J.; Shi, S.; Wang, X.; and Li, H. 2023. 3D object detection for autonomous driving: A comprehensive survey. *International Journal of Computer Vision*, 131(8): 1909–1963.
- Nabati, R.; and Qi, H. 2021. Centerfusion: Center-based radar and camera fusion for 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1527–1536.
- Park, J.; Xu, C.; Yang, S.; Keutzer, K.; Kitani, K. M.; Tomizuka, M.; and Zhan, W. 2022. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. In *The Eleventh International Conference on Learning Representations*.
- Philion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 194–210. Springer.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, 234–241. Springer.
- Singh, A. D.; Ba, Y.; Sarker, A.; Zhang, H.; Kadambi, A.; Soatto, S.; Srivastava, M.; and Wong, A. 2023. Depth estimation from camera image and mmwave radar point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9275–9285.
- Song, Z.; Yang, L.; Xu, S.; Liu, L.; Xu, D.; Jia, C.; Jia, F.; and Wang, L. 2024. Graphbev: Towards robust bev feature alignment for multi-modal 3d object detection. *arXiv preprint arXiv:2403.11848*.
- Svenningsson, P.; Fioranelli, F.; and Yarovoy, A. 2021. Radar-pointgcn: Graph based object recognition for unstructured radar point-cloud data. In *2021 IEEE Radar Conference (RadarConf21)*, 1–6. IEEE.
- Ulrich, M.; Braun, S.; Köhler, D.; Niederlöhner, D.; Faion, F.; Gläser, C.; and Blume, H. 2022. Improved Orientation Estimation and Detection with Hybrid Object Detection Networks for Automotive Radar. In *arXiv preprint arXiv:2205.02111*.
- Wang, S.; Liu, Y.; Wang, T.; Li, Y.; and Zhang, X. 2023. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3621–3631.
- Wang, Y.; Guizilini, V. C.; Zhang, T.; Wang, Y.; Zhao, H.; and Solomon, J. 2022a. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, 180–191. PMLR.
- Wang, Y.; Zhang, X.; Yang, T.; and Sun, J. 2022b. Anchor detr: Query design for transformer-based detector. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 2567–2575.
- Wu, Z.; Chen, G.; Gan, Y.; Wang, L.; and Pu, J. 2023. Mv-fusion: Multi-view 3d object detection with semantic-aligned radar and camera fusion. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2766–2773. IEEE.
- Xiong, K.; Gong, S.; Ye, X.; Tan, X.; Wan, J.; Ding, E.; Wang, J.; and Bai, X. 2023. Cape: Camera view position embedding for multi-view 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21570–21579.
- Yan, J.; Liu, Y.; Sun, J.; Jia, F.; Li, S.; Wang, T.; and Zhang, X. 2023. Cross modal transformer: Towards fast and robust 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18268–18278.
- Yan, Y.; Mao, Y.; and Li, B. 2018. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10): 3337.
- Yang, B.; Guo, R.; Liang, M.; Casas, S.; and Urtasun, R. 2020. Radarnet: Exploiting radar for robust perception of dynamic objects. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, 496–512. Springer.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.
- Zheng, L.; Li, S.; Tan, B.; Yang, L.; Chen, S.; Huang, L.; Bai, J.; Zhu, X.; and Ma, Z. 2023. Rcfusion: Fusing 4-d radar and camera with bird’s-eye view features for 3-d object detection. *IEEE Transactions on Instrumentation and Measurement*, 72: 1–14.
- Zhou, D.; Kang, B.; Jin, X.; Yang, L.; Lian, X.; Jiang, Z.; Hou, Q.; and Feng, J. 2021. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*.
- Zhou, T.; Chen, J.; Shi, Y.; Jiang, K.; Yang, M.; and Yang, D. 2023. Bridging the view disparity between radar and camera features for multi-modal fusion 3d object detection. *IEEE Transactions on Intelligent Vehicles*, 8(2): 1523–1535.