

RCM-Fusion: Radar-Camera Multi-Level Fusion for 3D Object Detection

Jisong Kim^{1,†}, Minjae Seong^{2,†}, Geonho Bang², Dongsuk Kum³ and Jun Won Choi^{4,*}

Abstract—While LiDAR sensors have been successfully applied to 3D object detection, the affordability of radar and camera sensors has led to a growing interest in fusing radars and cameras for 3D object detection. However, previous radar-camera fusion models could not fully utilize the potential of radar information. In this paper, we propose *Radar-Camera Multi-level fusion (RCM-Fusion)*, which attempts to fuse both modalities at feature and instance levels. For feature-level fusion, we propose a *Radar Guided BEV Encoder* which transforms camera features into precise BEV representations using the guidance of radar Bird’s-Eye-View (BEV) features and combines the radar and camera BEV features. For instance-level fusion, we propose a *Radar Grid Point Refinement module* that reduces localization error by accounting for the characteristics of the radar point clouds. The experiments on the public nuScenes dataset demonstrate that our proposed RCM-Fusion achieves state-of-the-art performances among single frame-based radar-camera fusion methods in the nuScenes 3D object detection benchmark. The code will be made publicly available.

I. INTRODUCTION

Accurate perception of the surrounding environments is crucial for successfully implementing self-driving robots or vehicles. Over the last decade, research into 3D object detection has focused on leveraging LiDAR’s precise spatial information and camera’s rich semantic information. However, LiDAR sensors may render them impractical due to their high cost, so radar sensors are emerging as a low-cost alternative. Radar sensors use radio waves to detect, locate and track objects in their vicinity. Radar sensors furnish both the 3D location of an object and its attributes, including Radar Cross Section (RCS) and Doppler velocity information. The 77GHz radars used for automotive applications have a longer wavelength than the LiDAR, making them more reliable in adverse weather conditions. Utilizing the advantages of these radar sensors, the previous literature has demonstrated that radar-camera fusion models can yield performance comparable to LiDAR-based 3D object detectors.

Recent advancements in radar-camera fusion have led to notable improvements in 3D object detection. These methods are roughly categorized into two distinct paradigms:

¹Department of Electrical Engineering, Hanyang University, Seoul, 04763, South Korea. jskim@spa.hanyang.ac.kr

²Department of Artificial Intelligence, Hanyang University, Seoul, 04763, South Korea. {mjseong, gbang}@spa.hanyang.ac.kr

³ Graduate School of Mobility, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 34141, South Korea. dskum@kaist.ac.kr

⁴Department of Electrical and Computer Engineering, College of Liberal Studies, Seoul National University, Seoul, 08826, Korea. junwchoi@snu.ac.kr

† Equally contributed authors

* Corresponding author

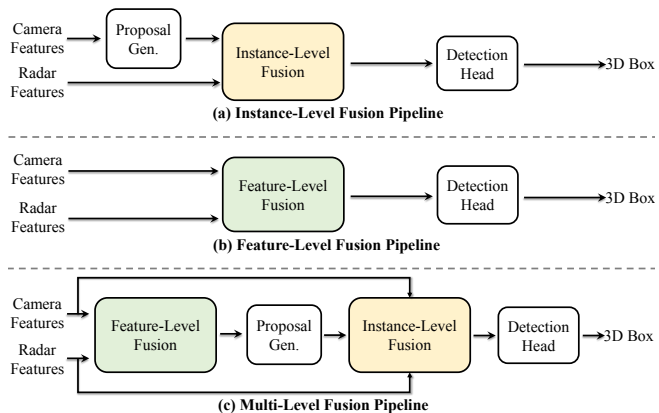


Fig. 1. **Radar-camera fusion pipeline:** Previous radar-camera fusion methods have primarily focused on fusion at single levels, such as (a) or (b). These approaches inherently suffer from limitations due to this level-specific fusion approach. To address this issue, we introduce a novel multi-level fusion method (c). By integrating multimodal information at multiple levels, this novel approach effectively overcomes the limitations of previous single-level fusion methods.

instance-level fusion [8], [13], [17] and feature-level fusion [9]. The instance-level fusion methods produce 3D box proposals from the camera features and associate the radar features represented in point clouds with 3D proposals in a camera-view to refine the initial results (see Fig. 1(a)). Despite their strengths, instance-level methods have a notable limitation: they do not fully exploit the radar’s capabilities during the initial proposal generation stage. Recently, feature-level fusion method, transformed the camera features into a bird’s eye view (BEV) domain through Lift, Splat, and Shoot (LSS) module [18] which was then fused with the radar BEV features (see Fig. 1(b)). Relying solely on either feature-level or instance-level fusion is limited in leveraging radar point information. To address this, we introduce a multi-level fusion framework that exploits radar and camera data at both the feature and instance levels (see Fig. 1(c)).

In this paper, we propose a new radar-camera fusion architecture, referred to as *Radar-Camera Multi-level Fusion (RCM-Fusion)* for 3D object detection. The proposed method achieves the fusion of radar and camera features at two different levels. RCM-Fusion first attempts feature-level fusion to produce a dense BEV feature map with which 3D proposals are generated. Then, it performs the instance-level fusion to refine the proposals through a grid point-based proposal feature fusion.

We devise Radar Guided BEV Query for feature-level fusion, which utilizes radar positional information to transform

image features into distinct BEV features. Then, we employ the Radar-Camera Gating module to obtain the weighted aggregation of multi-modal BEV feature maps. Such an adaptive feature aggregation module is integrated into transformer layers, which successively decode the dense BEV query features. For instance-level fusion, we propose the Proposal-aware Radar Attention module to obtain radar point features that consider their relevance to the 3D proposal. Then, we employ a novel adaptive grid point pooling method, Radar Grid Point Pooling, which accounts for the positional uncertainty of radar points for instance-level fusion. Finally, to improve the efficacy of radar data utilization and obtain high-quality radar points for RCM-Fusion, we modified the basic radar data filtering and multi-sweep methods used in the nuScenes dataset [1]. As a result, our model significantly outperformed camera-only baseline models and achieved state-of-the-art performance among the radar-camera fusion methods on the nuScenes 3D object detection benchmark [1].

The key contributions of our work are as follows:

- We present a state-of-the-art radar-camera fusion method for 3D object detection. The proposed RCM-Fusion combines radar and camera features at both feature and instance levels. This is in contrast to the existing radar-camera fusion methods [8], [9], [13], [17], which perform specific-level fusion only.
- We propose a **feature-level fusion** which transforms the image features into accurate BEV features through Radar Guided BEV Query and adaptively fuses multi-modal BEV features via Radar-Camera Gating modules. This feature-level fusion generates a dense BEV feature map, which is used to produce the initial 3D proposals.
- We propose an **instance-level fusion** that employs Proposal-aware Radar Attention to reduce the impact of irrelevant points and generates features which refine the initial results through Radar Grid Point Pooling.
- RCM-Fusion records a state-of-the-art performance of **50.6%** mAP and **58.7%** NDS in single frame-based radar-camera fusion methods on the nuScenes test set.

II. RELATED WORK

A. Radar-camera sensor fusion for 3D object detection

Radar-camera fusion models have been conducted in order to leverage radar positional information to supplement the camera’s limited depth information. The challenge of radar-camera fusion is to devise a strategy that effectively integrates data from various viewpoints. GriFNet [7] utilizes 3D anchors and 3D proposals to effectively extract features from radar and camera data. CramNet [5], meanwhile, leverages a 2D segmentation model to define the foreground section of both features and then fuses them in a unified 3D environment. CenterFusion [17], CRAFT [8] reduce localization and velocity errors by employing radar information based on

camera-based 3D object detection results. RADIANT [13] designed a network to estimate the positional offset between the radar returns and the object center in order to address the radar-camera association problem. CRN [9] utilizes camera BEV features, which are transformed from perspective view camera features using radar assistance, and fuses them with radar BEV features for 3D object detection.

Our RCM-Fusion differs from the aforementioned methods in that it performs radar-camera fusion at both feature and instance levels. Additionally, while the recent feature-level fusion method, CRN [9], necessitates depth ground truth labels for training its depth prediction network, our approach does not have this requirement.

B. Grid-based refinement for two-stage 3D object detection

Two-stage 3D object detectors [3], [16], [20], [23] using LiDAR data have been actively developed to refine initial prediction results by leveraging LiDAR point clouds within proposal boxes. In [3], [16], [20], virtual grid points were used to refine initial results. While PV-RCNN [20] and Voxel R-CNN [3] used uniformly spaced virtual grid points for refinement, Pyramid R-CNN [16] proposed a method that can reliably refine even when the internal points in proposals are sparse and irregularly distributed.

Our RCM-Fusion uses a novel grid-based refinement approach, which adaptively constructs 2D virtual grids, taking into account the directional uncertainty associated with radar measurements.

III. PROPOSED METHOD

We propose a multi-level fusion approach to efficiently fuse radar point clouds and surrounding images. Fig. 2 illustrates the overall architecture of our model. We consider a configuration where multi-view camera images and radar point clouds are used for object detection. Radar points and camera images are passed through separate backbone networks to generate feature maps for each sensor. The Radar-Guided BEV Encoder efficiently converts the camera view features to BEV and fuses them with the radar BEV features. Then, initial 3D proposals are obtained from the fused BEV feature maps using the transformer-based detection head. Finally, the Radar Grid Point Refinement module is used to refine the initial results, maximize the utilization of the radar data, and improve object detection accuracy.

A. Radar and Camera Backbones

Prior to combining radar and camera features, features are extracted using separate backbone networks for radar and camera modalities. Multi-scale image features F_C are generated from each of multi-view camera images using a ResNet backbone with FPN [11], which are commonly employed in existing camera-based 3D object detection models. For radar data, a PillarNet backbone [19] is used to obtain the radar feature map in BEV domain. The preprocessed radar points represented by 3D coordinates (x, y, z) , RCS, and Doppler velocity (v_x, v_y) are encoded by the 2D sparse convolution and passed through 2D convolution layers to obtain the radar BEV feature map $F_R \in \mathbb{R}^{H \times W \times C}$.

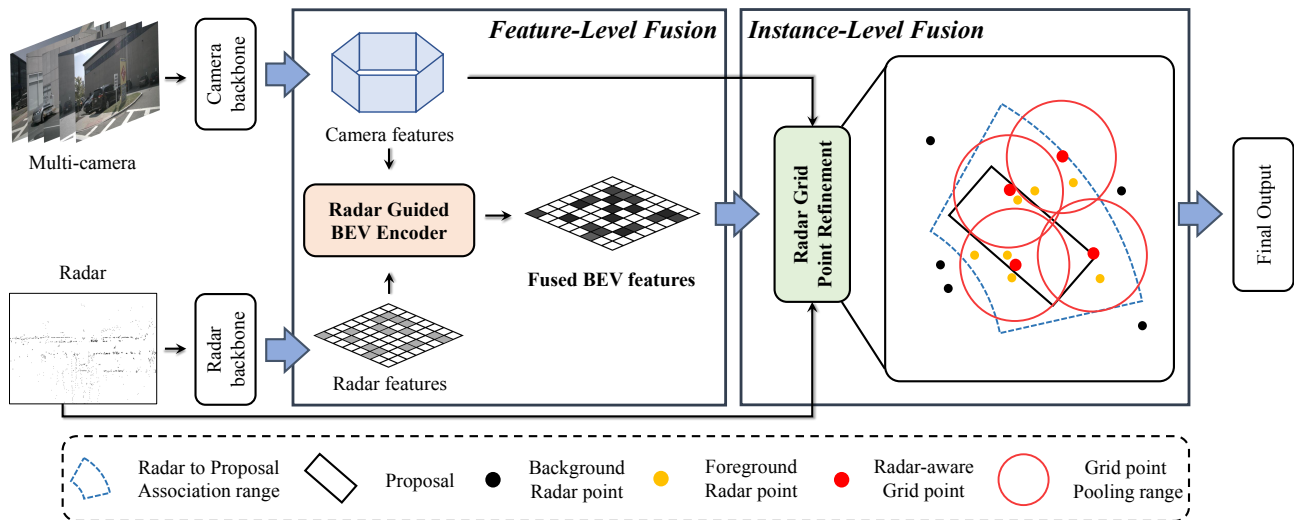


Fig. 2. **Overall architecture of the proposed RCM-Fusion:** RCM-Fusion uses each backbone network to obtain the radar features and camera features. The multi-modal features are then fused in the Radar Guided BEV Encoder module for feature-level fusion. For instance-level fusion, the Radar Grid Point Refinement module refines the initial results with a novel grid feature pooling method.

B. Radar Guided BEV Encoder

The Radar Guided BEV Encoder utilizes F_R and F_C as inputs to generate enhanced BEV features. First, we use F_R to create a Radar Guided BEV Query (RGBQ) containing positional information from the radar. Then, we use the RGBQs to transform F_C features into enhanced BEV features. Finally, the Radar-Camera Gating (RCG) performs the gated aggregation of enhanced BEV and encoded F_R features taking into account the difference in the amount of information between the two BEV features.

Radar Guided BEV Query. In order to achieve feature-level fusion of the radar and camera data, we need to transform F_C to a BEV representation. Our baseline model, BEVFormer [10], leverages the BEV query concept and transformer-based structure to perform view transformation. To further boost the accuracy of this view transformation, depth information in radar data can be utilized. Towards this goal, we use the positional information provided by the radar BEV feature map F_R to obtain an enhanced BEV query termed the Radar Guided BEV query $Q^{RG} \in \mathbb{R}^{H \times W \times C}$. This Radar Guided BEV query guides the camera-view to BEV feature transformation. Specifically, the Radar Guided BEV query Q^{RG} is generated by concatenating F_R with the learnable BEV queries $Q \in \mathbb{R}^{H \times W \times C}$ and passing it through a deformable self-attention (DeformAttn) module [25] as

$$Q_p^{RG} = \sum_{V \in \{Q, F_R\}} \text{DeformAttn}(Q_p, p, V), \quad (1)$$

where Q_p^{RG} and Q_p indicate the queries located at BEV plane pixel $p = (x, y)$. Then, Q_p^{RG} is used to decode the camera features through spatial-cross attention (SCA) block [10], generating the refined camera BEV features B_C as

$$B_C = \text{SCA}(Q_p^{RG}, F_C), \quad (2)$$

where SCA block is an operation that projects Q_p^{RG} to image features and then performs deformable cross-attention [25].

Radar-Camera Gating. Refined BEV feature map B_C and encoded F_R are fused by their weighted combination

$$F'_R = \text{MLP}_0(F_R), \quad (3)$$

$$B_{RC} = \{\sigma(\text{Conv}_C[B_C \oplus F'_R]) \odot B_C\} \oplus \{\sigma(\text{Conv}_R[F'_R \oplus B_C]) \odot F'_R\}, \quad (4)$$

where B_{RC} denotes the fused BEV feature map, F'_R denotes encoded radar BEV feature map, $\text{MLP}(\cdot)$ denotes the channel-wise MLPs, $\sigma(\cdot)$ denotes a sigmoid function, \odot and \oplus denote the element-wise multiplication and summation, respectively. Conv_C and Conv_R are convolutional layers for camera and radar, respectively. The gating operation adaptively weights the F_R and B_C taking into account information from both sides. This gating operation is implemented using a sigmoid function followed by convolutional layers.

Next, the fused BEV feature map B_{RC} is passed through normalization and feed-forward network similar to our baseline models, and the final BEV feature map is produced by repeating the BEV encoder section L times. When using the Radar Guided BEV Encoder, an enhanced feature map can be obtained based on radar positional information compared to the existing camera-only method. As seen in the activation maps in Figure 3, the feature map produced by our RCM-Fusion can resolve the region of the objects better than the baseline, BEVFormer-S [10], which does not use radar information for BEV transformation.

C. Radar Grid Point Refinement

In this section, we propose a novel instance-level fusion based on the Radar Grid Point Refinement module. The Proposal-aware Radar Attention (PRA) takes 3D proposals and their associated radar points as input and utilizes MLP-based attention layers to determine the importance of each radar point. Then, Radar Grid Point Pooling (RGPP) considers the characteristics and distribution of radar points to




Fig. 3. **Visualization of the BEV feature map:** Left: BEV feature map from BEVFormer-S [10], a single-frame version of BEVFormer. Right: BEV feature map obtained from Radar Guided BEV Encoder. Compared to the left figure, the right figure shows that the positional information provided by the radar BEV feature map allows the features to better localize the regions where objects exist.

sample the virtual grid points, and integrates the radar point features and multi-scale image features into the grid points for refinement. Finally, the refined features and the initial proposal features are combined to produce the final output.

Proposal-aware Radar Attention. First, we associate the radar points with 3D proposals using Soft Polar Association (SPA) method [8]. The SPA method transforms 3D proposals and radar points into polar coordinates, and maps radar points within a fixed azimuth and radial range to the 3D proposal. As the range for associating radar points expands, more radar points can contribute to the refinement. However, some of these points might not be relevant to the 3D proposal. PRA module can be used to solve this problem. Let b denote one of the N proposals found by the first stage detection, where $b = (\mathbf{c}, w, l, h, \theta, \mathbf{v}_{pred})$ consists of 3D proposal center location \mathbf{c} , 3D dimension (w, l, h) , orientation θ and velocity \mathbf{v}_{pred} . The K radar points associated with b are denoted by $\{r_k\}_{k=1}^K$, and the 3D position of the k th point r_k is $\mathbf{u}_k \in \mathbb{R}^3$. We encode the radar point features and relative positions between \mathbf{c} and \mathbf{u}_k to calculate the importance weight s_k for the k th point and obtain the attended radar point features a_k as

$$s_k = \text{MLP}_2([\text{MLP}_1(r_k); \delta(\mathbf{c} - \mathbf{u}_k)]) \quad (5)$$

$$a_k = \text{softmax}(\{s_k\}_{k=1}^K)_k \odot \text{MLP}_3(r_k), \quad (6)$$

where $\forall k \in \{1, \dots, K\}$

where $\delta(\cdot)$ denotes the positional encoding and $\text{softmax}(\cdot)$ denotes the Softmax function.

Radar Grid Point Pooling. The positioning and quantity of grid points have the most significant impact on the efficacy of the grid point-based refinement module [3], [16], [20]. In particular, radar points have a very low point density and a high positional error compared to LiDAR, which makes the grid point sampling more critical. Therefore, we propose the RGPP module, an adaptive grid point pooling technique that takes into account the characteristics of radar points. As shown in Fig. 4, the velocity vector \mathbf{v}_{pred} of 3D proposals, can be decomposed as the tangential velocity \mathbf{v}_{tan} and the radial velocity \mathbf{v}_{rad} . For the k^{th} radar point r_k , T grid points $\{g_k^t\}_{t=0}^{T-1}$ are generated around the radar point's position \mathbf{u}_k

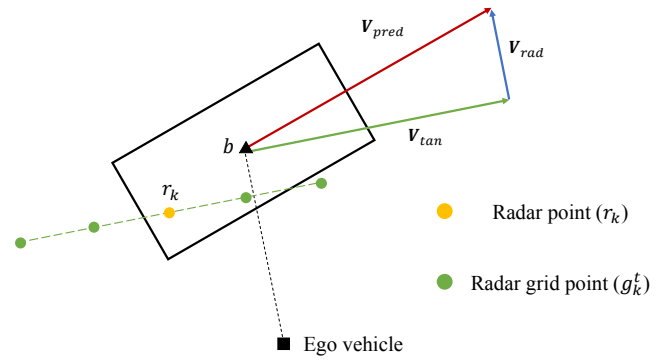


Fig. 4. **Illustration of novel adaptive grid point sampling:** For each radar point associated with the proposal, T grid points are generated in a tangential direction.

according to

$$\gamma = \begin{cases} \rho_{min}, & |\mathbf{v}_{tan}| \leq \rho_{min} \\ |\mathbf{v}_{tan}|, & \rho_{min} < |\mathbf{v}_{tan}| < \rho_{max} \\ \rho_{max}, & |\mathbf{v}_{tan}| \geq \rho_{max} \end{cases} \quad (7)$$

$$g_k^t = \gamma \cdot \left(\frac{t}{T-1} - \frac{1}{2} \right) \cdot \frac{\mathbf{v}_{tan}}{|\mathbf{v}_{tan}|} + \mathbf{u}_k, t = 0, \dots, T-1. \quad (8)$$

As shown in Fig. 4, we create the grid points in the direction of \mathbf{v}_{tan} . This arrangement is because radar points tend to be noisier in the tangential direction, so the grid points need to be located in this direction to accommodate the best radar point features. The spacing between the adjacent grid points is determined by γ , which is proportional to the magnitude of \mathbf{v}_{tan} . Since \mathbf{v}_{tan} have a predicted value, it may not be accurate, and grid points may not be generated correctly. Therefore, we set the upper bound ρ_{max} and the lower bound ρ_{min} for the range of γ . For each 3D proposal, KT grid points are generated from K radar points, and then Farthest Point Sampling (FPS) is utilized to select M grid points $\{g_m\}_{m=1}^M$ to maintain a fixed number of grid points.

Next, we use the grid points $\{g_m\}_{m=1}^M$ as a basis for extracting radar point features and image features and generating the fused grid features. First, the set abstraction (SetAbs) [20] encodes the radar points around each grid point g_m to produce the radar point features F_m^{pts}

$$F_m^{pts} = \text{SetAbs}(\{a_k\}_{k=1}^K, \{\mathbf{u}_k\}_{k=1}^K, g_m). \quad (9)$$

In parallel, the grid points are projected onto F_C in the camera view and then the camera features F_m^{img} are obtained through a bilinear interpolation

$$F_m^{img} = \text{Bilinear}(F_C, \text{proj}(g_m)), \quad (10)$$

where $\text{proj}(\cdot)$ denotes the projection matrix to the camera view. Finally, the proposal features F_m^{obj} are obtained by fusing F_m^{pts} and F_m^{img}

$$F_m^{obj} = \text{maxpool}(F_m^{pts} \oplus F_m^{img}), \quad (11)$$

and these features are fused with the initial proposal features to perform 3D proposal refinement.

TABLE I

PERFORMANCE COMPARISONS WITH SINGLE FRAME 3D OBJECT DETECTOR ON THE nuSCENES TEST SET. ‘C’, AND ‘R’ REPRESENT CAMERA, AND RADAR, RESPECTIVELY. * ARE TRAINED WITH EXTERNAL DATA.

Method	Input	Backbone	NDS	mAP	mATE	mASE	mAOE	mAVE	mAAE
BEVDet* [4]	C	V2-99	48.8	42.4	0.524	0.242	0.373	0.950	0.148
PETR [12]		R101	45.5	39.1	0.647	0.251	0.433	0.933	0.143
Graph-DETR3D [2]		R101	47.2	41.8	0.668	0.250	0.440	0.876	0.139
Ego3RT [15]		R101	44.3	38.9	0.599	0.268	0.470	1.169	0.172
BEVFormer-S [10]		R101	46.2	40.9	0.650	0.261	0.439	0.925	0.147
PolarFormer [6]		R101	47.0	41.5	0.657	0.263	0.405	0.911	0.139
M ² BEV [22]		X101	47.4	42.9	0.583	0.254	0.376	1.053	0.190
CenterFusion [17]	R+C	DLA34	44.9	32.6	0.631	0.261	0.516	0.614	0.115
CRAFT [8]		DLA34	52.3	41.1	0.467	0.268	0.456	0.519	0.114
RCM-Fusion-R50		R50	53.5	45.2	0.496	0.264	0.476	0.554	0.122
RCM-Fusion-R101		R101	58.7	50.6	0.465	0.254	0.384	0.438	0.121

TABLE II

ABLATION STUDY CONDUCTED ON nuSCENES VAL SET. RGBQ: THE RADAR GUIDED BEV QUERY. RCG: THE RADAR-CAMERA GATING. RGPP: THE RADAR GRID POINT POOLING. PRA: THE PROPOSAL-AWARE RADAR ATTENTION.

Feature Level		Instance Level		Performance			
w/ RGBQ	w/ RCG	w/ RGPP	w/ PRA	NDS	mAP	FPS	Params
				44.9	37.9	2.7	66.24M
✓				49.4	42.5	2.5	85.44M
✓	✓			51.6	43.2	2.4	87.41M
✓	✓	✓		52.3	43.7	2.3	87.60M
✓	✓	✓	✓	52.9	44.3	2.3	87.62M

TABLE III

COMPARISON OF GRID POINT SAMPLING METHODS. W/O GRID POINTS: REPLACES RADAR POINTS WITH GRID POINTS. FIXED GRID POINTS: GENERATES A SET NUMBER OF GRID POINTS [20].

Method	NDS	mAP
w/o Grid Points	52.4	43.8
Fixed Grid Points	52.5	44.0
Adaptive Grid Points	52.9	44.3

IV. EXPERIMENTS

A. Datasets

We evaluated our model on the nuScenes [1] dataset, a benchmark for autonomous driving. The nuScenes dataset consists of 700, 150, and 150 scenes for *train*, *val*, and *test*, respectively. The nuScenes dataset evaluates 3D object detection performance using mean Average Precision (mAP) and the nuScenes Detection score (NDS). The mAP is calculated from the prediction and ground truth’s BEV center distance, averaged over 0.5, 1, 2, and 4 meters thresholds. The nuScenes Detection Score (NDS) is calculated by combining mAP with five True Positive (TP) metrics: translation, scale, orientation, velocity, and attribute errors.

TABLE IV

PERFORMANCE ANALYSIS ON nuSCENES VAL SET.

Metric	# of Encoder Layers		
	1	2	3
NDS	51.2	52.9	52.7
mAP	42.3	44.3	44.1

(a) # of Encoder Layers L

Metric	# of Grid Point				Metric	# of Grid Point FPS			
	3	5	7	9		16	32	64	128
NDS	52.6	52.7	52.9	52.9	NDS	52.4	52.5	52.9	52.4
mAP	44.0	44.2	44.3	44.2	mAP	44.1	44.3	44.3	44.0

(b) # of Grid Points Per Radar Point, T (c) # of Selected Grid Points, M

B. Implementation Details

As a baseline model, we used a static version of BEVFormer [10], BEVFormer-S, that only utilized single-frame data. We used the pre-trained weights of FCOS3D [21] as a camera backbone. We refer to our method as RCM-Fusion-R101 when ResNet-101 is used. We also call RCM-Fusion-R50 when ResNet-50 is used. We used the image size of 900×1600 for RCM-Fusion-R101 and the size of 450×800 for RCM-Fusion-50. For radar, we accumulate six previous radar sweeps. We did not use pre-trained weights for the radar backbone and trained the entire model with random initialization in an end-to-end fashion. We applied the Image Data Augmentation and BEV Data Augmentation techniques in BEVDet [4]. We used the AdamW [14] optimizer with a learning rate of $2e-4$ and trained 24 epochs, employing the class balancing strategy CBGS proposed in [24]. All ablation studies used $1/7$ of the nuScenes *train* and applied the CBGS technique for RCM-Fusion-R101.

C. Results on the nuScenes Dataset

Table I presents the performance of our RCM-Fusion on the nuScenes *test* set. We include the performance of the existing camera-based methods and radar-camera fusion methods. For a fair comparison, we focus solely on single frame-based methods. Thus, CRN [9] was excluded because its single frame performance was not provided in the original paper. Compared with our baseline model, BEVFormer-S




Fig. 5. **Qualitative results:** We visualized the predictions of RCM-Fusion (green boxes) and camera-only baseline model, BEVFormer-S (blue boxes), alongside the GT (red boxes) on a bird's eye view. The figures demonstrate that the RCM-Fusion produces more accurate results (in the orange regions), and it can detect the objects the camera-only baseline fails to identify (in the yellow regions).

[10], RCM-Fusion-R101 achieves 12.5% and 9.7% performance gains in NDS and mAP, respectively. Note that RCM-Fusion-R101 outperforms all other camera-based and radar-camera fusion methods by significant margins. Despite the ResNet50 backbone being considered inferior to the DLA34 backbone, RCM-Fusion-R50 also outperforms CRAFT [8] by 1.2% in NDS and 4.1% in mAP.

D. Ablation Studies on the nuScenes Val Set

Component Analysis. Table II presents the ablation study that identifies the contributions of Radar Guided BEV Query (RGBQ), Radar-Camera Gating (RCG), Radar Grid Point Pooling (RGPP) and Proposal-aware Radar Attention (PRA) to the overall performance. Utilizing radar pillar features to refine BEV queries, RGBQ improves the performance of the baseline by 4.6% in mAP and 4.5% in NDS. RCG that adaptively fuses multi-modal BEV features offers additional performance gains of 0.7% in mAP and 2.2% in NDS. In the instance-level fusion, RGPP, which fuses the multi-modal features based on the adaptive grid points, further improves the performance by 0.5% in mAP and 0.7% in NDS. Finally, adding the PRA module leads to a 0.6% performance gain in both mAP and NDS. Note that the instance-level fusion results in only a slight increase in both frames per second (FPS) and the number of parameters while it achieves an 1.3% gain in NDS and 1.0% gain in mAP. These results demonstrate the effectiveness of our instance-level fusion.

Qualitative Results. In Fig. 5, the detection results from RCM-Fusion are visualized for various scenarios. By utilizing the range information provided by the radar, RCM-Fusion can localize the objects more accurately than the camera-only baseline model, BEVFormer-S. Moreover, our model can detect some objects that the camera-only baseline model fails to identify.

Effects of Radar Grid Point Sampling. We investigate the effectiveness of the proposed grid point generation method. Table III compares the proposed method with other grid generation methods. The w/o Grid Points (w/o GP) is a method that does not generate grid points and directly uses the radar points associated with the proposal for refinement. The Fixed Grid Point method is a method that generates a fixed number of grid points regularly distributed within a

proposal box, similar to [20]. While the Fixed Grid Point method was effective for LiDAR data, its performance is not satisfactory with radar data because of its sparse distribution and increased errors along the horizontal axis. We note that the Fixed Grid Point method barely improves the w/o GP method. However, by employing Adaptive Grid Points that adjust both direction and spacing based on the distribution of radar points, our grid point sampling method achieved better performance over the baselines.

Performance versus Hyper-parameters. Table IV presents the performance of RCM-Fusion in relation to several hyper-parameters. Specifically, Table IV (a) investigates the performance as a function of the number of layers in the Radar Guided BEV encoder. Our experimental results indicate that a performance peak is achieved when the number of layers is set to two. Table IV (b) presents the performance versus the number of grid points T generated per each radar point. We found that the best value for T is 7. Table IV (c) shows the best value for the number of grid points M used in the FPS algorithm is 64.

V. CONCLUSIONS

In this paper, we proposed a novel radar-camera 3D detector, RCM-Fusion. RCM-Fusion fused the radar and camera data at both the feature-level and instance-level for enhanced 3D object detection. We demonstrated that the Radar Guided BEV Encoder could achieve effective transformation of camera features into BEV, leveraging the radar features. Moreover, the Radar Grid Point Refinement proposed a novel grid pooling technique, which fused both radar and camera features based on the adaptive grid points to refine the initial 3D proposals. For future work, we plan to expand our approach towards designing a time-efficient multi-frame radar-camera fusion method.

VI. ACKNOWLEDGMENT

This work was supported by Institute for Information & communications Technology Promotion(IITP) and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2021-0-01314,Development of driving environment data stitching technology to provide data on shaded areas for autonomous vehicles and 2020R1A2C2012146)

REFERENCES

- [1] 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 multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11621–11631, 2020.
- [2] Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qinhong Jiang, and Feng Zhao. Graph-detr3d: rethinking overlapping regions for multi-view 3d object detection. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5999–6008, 2022.
- [3] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. *arXiv preprint arXiv:2012.15712*, 1(2):4, 2020.
- [4] 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.
- [5] Jyh-Jing Hwang, Henrik Kretzschmar, Joshua Manela, Sean Rafferty, Nicholas Armstrong-Crews, Tiffany Chen, and Dragomir Anguelov. Cramnet: Camera-radar fusion with ray-constrained cross-attention for robust 3d object detection. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVIII*, pages 388–405. Springer, 2022.
- [6] Yanqin Jiang, Li Zhang, Zhenwei Miao, Xiatian Zhu, Jin Gao, Weiming Hu, and Yu-Gang Jiang. Polarformer: Multi-camera 3d object detection with polar transformers. *arXiv preprint arXiv:2206.15398*, 2022.
- [7] Youngseok Kim, Jun Won Choi, and Dongsuk Kum. Grif net: 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.
- [8] Youngseok Kim, Sanmin Kim, Jun Won Choi, and Dongsuk Kum. Craft: Camera-radar 3d object detection with spatio-contextual fusion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.
- [9] Youngseok Kim, Sanmin Kim, Juyeb Shin, Jun Won Choi, and Dongsuk Kum. Crn: Camera radar net for accurate, robust, efficient 3d perception. *arXiv preprint arXiv:2304.00670*, 2023.
- [10] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, pages 1–18. Springer, 2022.
- [11] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017.
- [12] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVII*, pages 531–548. Springer, 2022.
- [13] Yunfei Long, Abhinav Kumar, Daniel Morris, Xiaoming Liu, Marcos Castro, and Punarjay Chakravarty. Radiant: Radar-image association network for 3d object detection. *arXiv preprint arXiv:2206.15398*, 2023.
- [14] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [15] Jiachen Lu, Zheyuan Zhou, Xiatian Zhu, Hang Xu, and Li Zhang. Learning ego 3d representation as ray tracing. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 129–144. Springer, 2022.
- [16] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid r-cnn: Towards better performance and adaptability for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2723–2732, 2021.
- [17] 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.
- [18] Jonah Philion and Sanja Fidler. 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*, pages 194–210. Springer, 2020.
- [19] Guangsheng Shi, Ruifeng Li, and Chao Ma. Pillarnet: Real-time and high-performance pillar-based 3d object detection. In *European Conference on Computer Vision*, pages 35–52. Springer, 2022.
- [20] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10529–10538, 2020.
- [21] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 913–922, 2021.
- [22] Enze Xie, Zhiding Yu, Daquan Zhou, Jonah Philion, Anima Anandkumar, Sanja Fidler, Ping Luo, and Jose M Alvarez. M² 2bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation. *arXiv preprint arXiv:2204.05088*, 2022.
- [23] Honghui Yang, Zili Liu, Xiaopei Wu, Wenxiao Wang, Wei Qian, Xiaofei He, and Deng Cai. Graph r-cnn: Towards accurate 3d object detection with semantic-decorated local graph. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*, pages 662–679. Springer, 2022.
- [24] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019.
- [25] Xizhou Zhu, Weiye Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.