

MVFusion: Multi-View 3D Object Detection with Semantic-aligned Radar and Camera Fusion

Zizhang Wu¹, Guilian Chen¹, Yuanzhu Gan¹, Lei Wang¹, Jian Pu²

Abstract—Multi-view radar-camera fused 3D object detection provides a farther detection range and more helpful features for autonomous driving, especially under adverse weather. The current radar-camera fusion methods deliver kinds of designs to fuse radar information with camera data. However, these fusion approaches usually adopt the straightforward concatenation operation between multi-modal features, which ignores the semantic alignment with radar features and sufficient correlations across modals. In this paper, we present MVFusion, a novel Multi-View radar-camera Fusion method to achieve semantic-aligned radar features and enhance the cross-modal information interaction. To achieve so, we inject the semantic alignment into the radar features via the semantic-aligned radar encoder (SARE) to produce image-guided radar features. Then, we propose the radar-guided fusion transformer (RGFT) to fuse our radar and image features to strengthen the two modals' correlation from the global scope via the cross-attention mechanism. Extensive experiments show that MVFusion achieves state-of-the-art performance (51.7% NDS and 45.3% mAP) on the nuScenes dataset. We shall release our code and trained networks upon publication.

I. INTRODUCTION

Significant progress of deep learning in recent years has achieved remarkable performance in the object detection task of autonomous driving [1], [2]. To realize a more robust perception for driving assistance, the cooperation of multiple sensor modalities, such as camera, LiDAR, and millimeter wave radar (radar), has attracted wide attentions in many research works [3]–[6].

Each sensor modality has its own merits [6]. Camera-based images provide high-resolution visual signals, which perform similar human visual perception, and even overcomes human limitation when equipped with multiple surround-view cameras. LiDAR data supplements depth information compared to the single camera perception, but spends more cost [6], [7]. However, the camera and LiDAR fail to offer robust perception under adverse weather, such as heavy rain, fog, and other low visibility conditions. In contrast, radar reveals more stability to adverse conditions with lower cost [8]. In addition, radar provides a farther detection range and more helpful features, such as velocity, radar cross section (RCS), distance, etc.

Nevertheless, radar points perform sparse and noisy, lacking adequate vertical measurement, so researchers usually fuse the radar knowledge with LiDAR or camera data [6], [9], [10], where the radar-camera fusion reveals the mainstream strategy. For example, CenterFusion [7] achieves center-based radar and camera fusion via the specific two-stage

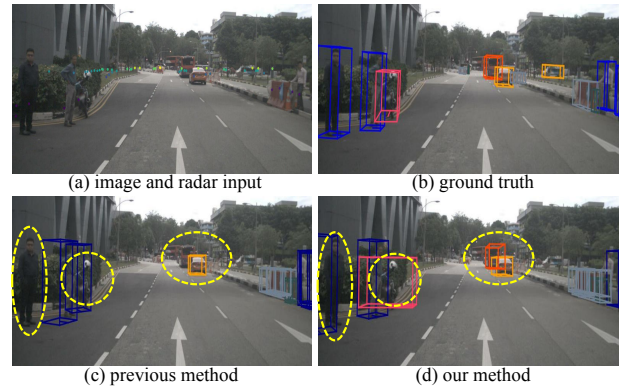


Fig. 1. Detection comparison between the camera-based approach [13] and our MVFusion. (a) The image and radar input, where the color of radar points denotes the distance to the radar. (b) The 3D detection ground truth. (c) The results of the camera-based approach [13], which misses the detection of the distant car and the near pedestrian. (d) Our method takes advantage of the semantic-aligned radar information to perform sufficient radar-camera fusion and successfully detects the missing car and pedestrian.

design. However, these fusion approaches usually adopt the straightforward concatenation operation between multi-modal features, which ignores the spatial connections and semantic alignment [7], [10], [11].

Moreover, the previous works [7], [12] focus on radar-camera fusion within monocular image content, there are few studies to survey the merging between multi-camera images and radar data. Multi-view camera-based 3D object detection [13], [14] captures visual information across different camera views, and conducts more excellent performance than the monocular camera-based detection. Yet the multi-view radar-camera fusion remains under exploration, where radar data could assist and promote better surrounding perception.

In this paper, we present a novel multi-view radar-camera fusion method to achieve semantic-aligned radar features and enhance the cross-modal information interaction from the global scope. As shown in Fig.1, our method catches more helpful information via the semantic alignment radar features and radar-camera fusion, to receive more correct detection.

In detail, considering the lack of corresponding semantic information of radar points, firstly we propose the semantic-aligned radar encoder (SARE) to extract image-guided radar features, which adopts the semantic indicator to align radar inputs and select the image-guided radar transformer to form strong radar representations.

After radar feature extraction, we propose the radar-guided fusion transformer (RGFT) to integrate the enhanced radar features with high-level image features. We utilize the cross-attention transformer to strengthen the cross-modal interac-

¹Zongmu Technology

²Fudan University

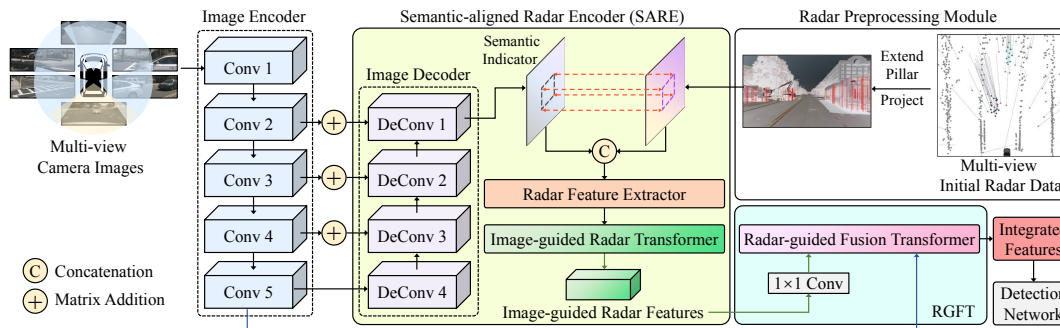


Fig. 2. The overview of our proposed **MVFusion**, which mainly consists of five components: the radar preprocessing module, the image encoder, the semantic-aligned radar encoder (SARE), the radar-guided fusion transformer (RGFT) and the detection network. SARE injects the semantic alignment into the radar features, and RGFT fuses our radar and image features, aiming to sufficiently promote the two modal’s interaction from the global scope. The multi-view radar representation refers to [15].

tion from global scope, which performs beneficial for the subsequent 3D detection.

Our main contributions are as follows:

(1) We are the first to explore the radar-camera multi-view fusion for 3D object detection and propose our **Multi-View radar-camera Fusion** approach: **MVFusion**, which utilizes visual semantics to obtain semantic-aligned radar features and applies robust fusion transformer to enhance the cross-modal information interaction.

(2) We propose the semantic-aligned radar encoder (**SARE**) to extract image-guided radar features, which adopts the semantic indicator to align radar data and image-guided radar transformer to produce strong radar representations. Moreover, we present a radar-guided fusion transformer (**RGFT**) to integrate the enhanced radar features with high-level image features, to promote the two modal’s sufficient correlation from the global scope via the cross-attention mechanism.

(3) Experiments show that **MVFusion** achieves state-of-the-art performance (51.7% NDS and 45.3% mAP) of the single-frame multi-view radar-camera fused 3D object detection on the standard nuScenes benchmark dataset.

II. RELATED WORK

A. Single-modality 3D Object Detection

Three-dimensional (3D) object detection reveals critical for various applications [9], [16]–[21], such as autonomous driving, etc. Previous works commonly utilize kinds of sensors (LiDAR, camera, and radar) to perform 3D object detection. LiDAR-based methods [9], [21]–[23] have achieved the most precise 3D measurements but the most expensive price among the three sensors. Then cameras achieve more affordable cost and provide abundant semantic information, though lacking depth information, which also attract extensive attention [1], [24]–[26], such as monocular 3D object detection [19], [20], [27] and multi-view 3D object detection [13], [14], [25], [28]. Specifically, the multi-view 3D object detection integrates information across different camera views, for producing a more robust surrounding scene understanding. BEVDet [14] lifts the multi-view visual features to frustum features, then fuses and splats the frustum features into BEV (Bird’s Eye View) space for detection.

DETR3D [25] obtains 3D object queries from 2D features of multiple camera images and generates reference points from these queries. PETR [13] transforms the camera frustum space into 3D coordinates, to generate 3D position-aware features for the subsequent transformer decoder. At last, radar provides important features like velocity, etc. in most weather conditions but achieves noisy position values, especially the height information [8]. Researchers usually adopt the radar data together with other sensors’ information to the multi-sensor fusion [8], [10], [12], [29]–[31] detection.

B. Fusion-based 3D Object Detection

The 3D object detection based on multi-modal fusion collects single-modal advantages to promote the detection performance [6], [10], [29], [30]. Many works [32]–[37] focus on the fusion of camera and LiDAR data. For example, DeepFusion [35] projects LiDAR points to the image plane and applies Learnable Align to capture the correlations with the cross-attention mechanism. Pointaugmenting [37] utilizes a BEVbased fusion approach with cross-LiDAR-camera data augmentation. In addition, the immunity to extreme conditions and the characteristic cost-effective make radar considerable in developing autonomous driving with more robustness. Recent methods [7], [10]–[12] have demonstrated the effectiveness based on the early, middle, and late fusion of radar and camera data. The approach [10] focuses on object-level fusion via deep representation learning to explore the interaction and global reasoning of different sensor features. CenterFusion [7] detects objects’ image centers to associate radar detections and generates radar features, then concatenates image and radar features for detection heads. However, these fusion methods usually perform straightforward concatenation operations, ignoring the spatial connections and semantic alignment between radar and image features. Besides, there are rare explorations for multi-view camera radar fusion. So in this study, we aim to enhance the feature correlations and strengthen the spatial information interaction between radar and camera data from multiple views.

III. METHOD

A. Overall Architecture

We demonstrate the framework of our multi-view radar-camera fusion approach in Fig. 2, which mainly consists of




Fig. 3. The schema of radar feature extractor (RFE), which consists of residual feature convolution blocks for the sparse radar features.

five components: the radar preprocessing module, the image encoder, the semantic-aligned radar encoder (SARE), the radar-guided fusion transformer (RGFT) and the detection network. For the image inputs, given surround-view images from six views, we adopt the image encoder to extract the multi-scale image features, where we denote the image features of each stage as $F_{image}^s = \text{Conv}_s(I)$, $s = 1, 2, 3, 4, 5$. For the radar points inputs from multi-view radars, we first employ the radar preprocessing module similar to [8] to obtain the radar representations with the same shape as images. Then we apply our semantic-aligned radar encoder (SARE) to produce the image-guided enhanced radar features. Afterward, we utilize the radar-guided fusion transformer (RGFT) to conduct the camera-radar fusion in terms of high-level features from the global scope, to create the integrated features. Finally, the detection network chooses the cross-attention decoder and 3D detection head, to efficiently decodes the object queries and perform the final 3D predictions.

The radar preprocessing module exhibits essential, for the initial radar data reveals 3D points with additional properties, thus we need to transform them to a suitable formation to fuse the camera data (in this work, we project the radar points to the image space). Moreover, the radar data performs sparse and noisy, so its direct utilization will affect the detection performance [8]. Specifically, we obey two steps to preprocess the radar data. Firstly, to avoid the noisy points disturbing, we filter the radar points with the 3D bounding boxes of ground truth, following [8]. Besides, considering the incorrect or missing vertical measurement, we expand the points to the pillar with a three-meter height. Then we project the radar pillars to the corresponding image view with camera calibration, where we replace the back-view data with blank points, for the vacancy of radar at the back of the ego-vehicle. Furthermore, we reserve another four significant characters to compose the five-channel radar maps (I_{radar}), where each point carries five properties: the position mask with zero or one (indicating the existence of radar points at the corresponding position of pixels); the distance; the radar cross section (RCS): a physical quantity of the echo intensity produced by the illumination of radar waves); the two radial velocities along X and Y axis.

B. Semantic-aligned Radar Encoder (SARE)

Through the image encoder, we can obtain multi-scale image features with both low-level and high-level semantics. Different from previous fusion methods [7], [12], we firstly achieve semantic alignment between our five-channel radar maps after preprocessing and semantic image features before




Fig. 4. The overview of our image-guided radar transformer (IGRT). IGRT assigns learnable position encoding to the radar features to further enhance the spatial information through the multi-head self-attention mechanism.

feeding into the radar feature extractor, as shown in Fig. 2. Specifically, the semantic-aligned radar encoder contains three modules: the image decoder, the radar feature extractor and the image-guided radar transformer.

To achieve semantic alignment, firstly we require a robust semantic indicator to align and filter the sparse radar inputs. Here we produce our image indicator via all stages' visual features, where we upsample the high-level features through deconvolution blocks and complement them with original low-level features by skip-connections [38], to reach a more robust semantic indicator (SI). We denote the i -th output D_i of the image decoder as:

$$D_i = \begin{cases} \text{DeConv}(F_{image}^5), & i = 4 \\ \text{DeConv}(D_{i+1}) + F_{image}^{i+1}, & i = 1, 2, 3 \end{cases} \quad (1)$$

where $DeConv$ stands for the deconvolution block. At last, we conduct 1×1 convolution to reduce the channel to one and concatenate the image indicator to the radar inputs after preprocessing. We set the image indicator's channel as one, for we desire the indicator as the image foreground mask and utilize the foreground's semantics and relative location within the image space to achieve more robust alignment between semantic points and radar points. Concatenated under the high resolution, we succeed in accurately aligning radar points with the foreground semantics which are inferred according to the high-level visual features.

After obtaining the semantic-aligned radar inputs, we put forward the radar feature extractor (RFE), as shown in Fig. 3. In the first convolution layer, we adopt the 1×1 convolution with stride of 2, and a max-pooling operation with kernel size of 2 to down-sample the features. Furthermore, three sparse radar features extraction blocks with stride of 2 are employed to extract the final coarse radar features: $40 \times 100 \times 256$ ($H \times W \times C$). We denote the coarse radar features $F_{radar} = \text{RFE}(I_{indicator}, I_{radar})$.

However, the universal convolution to extract radar features [7], [10]–[12] performs the suboptimal performance, which ignores the implicit position and semantic relationships. Thus we propose the image-guided radar transformer (IGRT) to model the long-range dependency and relationships for the coarse radar features. In particular, firstly we flatten the features in H and W dimensions into the flattened sequences with the reshaping operation. Then we introduce the multi-head self-attention mechanism to distribute attention on different positions simultaneously and enhance our radar features, as shown in Fig. 4. The self-attention achieves the overall and intelligent interaction to acquire sufficient

correlations. Furthermore, we adopt the learnable position encoding [39], which reveals benefits for improving the spatial information learning of radar features. We describe the self-attention procedure as follows:

$$Q, K, V = (F_{radar} + E_{pos})\mathbf{W}_{Q,K,V} \quad (2)$$

$$\text{IGRT-Attn} = \text{Softmax}\left(\frac{QK^T}{\sqrt{C/h}}\right)V \quad (3)$$

$$F'_{radar} = \text{MLP}(\text{LN}(\text{IGRT-Attn})) \quad (4)$$

Q , K , and V denote the Query, Key, and Value. \mathbf{W}_Q , \mathbf{W}_K and \mathbf{W}_V are learnable parameters. E_{pos} stands for the learnable position encoding. h denotes the number of heads for the multi-head self-attention. LN means layer normalization and MLP means the multi-layer perceptron block.

C. Radar-guided Fusion Transformer (RGFT)

With the alignment and guidance of semantic information with image features, we attain enhanced and robust radar features. To provide better fusion features for the detection network, we propose the radar-guided fusion transformer for camera-radar fusion in terms of high-level features from the global scope. The earlier fusion approaches [6], [9], [10] usually adopt the straightforward concatenation operation between multi-modal features, which ignores the sufficient and global spatial correlations across modals. Notably, some transformer-based methods [13], [28] succeed in managing the cross-modal problem via the cross-attention mechanism within the decoder. Thus here we introduce the cross-attention mechanism to fuse our radar and image features, aiming to strengthen the two modal's interaction from the global field, as shown in Fig. 5.

Different from the above radar encoder, for the fusion transformer, we remove the multi-head attention module to avoid the cost of large computation. In addition, we utilize the 1×1 convolution (Fig. 2) to reduce the channel of high-level radar features before sending into the RGFT, and adopt another 1×1 convolution (Fig. 5) to expand the channel of F_{image}^5 to the fusion features' wanted channel.

$$F''_{radar} = \text{conv}_{1 \times 1}^a(F'_{radar}) \quad (5)$$

$$F''_{image} = \text{conv}_{1 \times 1}^b(F_{image}^5 + E'_{pos}) \quad (6)$$

Furthermore, we adopt the concatenation of the radar and image features to create the cross-attention module's query, for only the relatively sparse radar features result in a convergence problem with exploding gradients (setting (c) in Table V). On the other hand, we develop the keys and values from the image features. Therefore, the cross-attention-related equation of our proposed RGFT can be defined as:

$$Q = \text{Concat}(F''_{radar}, F''_{image})\mathbf{W}_Q \quad (7)$$

$$K = F''_{image}\mathbf{W}_K, \quad V = F''_{image}\mathbf{W}_V \quad (8)$$

$$\text{RGFT-Attn} = \text{Softmax}\left(\frac{QK^T}{\sqrt{C}}\right)V \quad (9)$$




Fig. 5. The overview of radar-guided fusion transformer (RGFT). Fusing the high-level radar and image features, RGFT achieves adequate correlation under the cross-attention mechanism.

where E'_{pos} is the learnable position encoding for the image features, and C denotes the channel number of F_{image}^5 . The final integrated features can be written as:

$$F_{integrated} = \text{MLP}(\text{LN}(\text{RGFT-Attn})) \quad (10)$$

With the radar-guided fusion transformer (RGFT) for adequate cross-modal correlations, we achieve robust integrated fusion features for the subsequent detection procedure.

D. Detection Network

We adopt the same multi-view decoder and 3D detection head as the previous work [13], where object queries interact with the integrated fusion features through the multi-head attention and feed-forward network. Finally, every object query predicts one 3D location and its category through the detection head, then we utilize the Hungarian algorithm [40] for label assignment between ground truths and predictions. We adopt the focal loss [41] for 3D classification and $L1$ loss for 3D regression:

$$L_{3D} = \lambda_{cls}L_{cls} + \lambda_{reg}L_{reg}. \quad (11)$$

IV. EXPERIMENTS

A. Datasets and Metrics

We validate the effectiveness of our method on the public large-scale dataset nuScenes [46], which contains a full 360-degree field of view provided by 6 cameras, 1 lidar and 5 radars. The dataset consists of 1000 driving scenes, with 700, 150 and 150 scenes for training, validation, and testing, respectively. We sample the corresponding sequences to frames at 2Hz and objects in each frame are annotated with 3D bounding boxes. We follow the official evaluation metrics of the dataset including nuScenes Detection Score (NDS) and mean Average Precision (mAP), along with mean Average Translation Error (mATE), mean Average Precision (mAP), mean Average Translation Error (mATE), mean Average Scale Error (mASE), mean Average Orientation Error (mAOE), mean Average Velocity Error (mAVE), mean Average Attribute Error (mAAE).

B. Implementation Details

Our method applies ResNet [47] and VoVNetV2 [43] as the feature extractor. We adopt the same multi-view decoder and 3D detection head as the previous work [13]. In addition, we also apply some basic experiment settings referred to [13]. Specifically, we choose 64 points along the depth axis following the linear-increasing discretization (LID) [48], and set the region to $[-61.2m, 61.2m]$ for the X and Y axis, and

TABLE I

COMPARISON OF SINGLE-FRAME STATE-OF-THE-ART WORKS ON THE nuSCENES TEST SET WITH DIFFERENT MODALS. † DENOTES ADOPTING THE DD3D [42] PRE-TRAINED V2-99 [43] BACKBONE.

Method	Modal	NDS \uparrow	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
InfoFocus [44]	L	0.395	0.395	0.363	0.265	1.132	1.000	0.395
PointPillars [9]	L	0.453	0.305	0.517	0.290	0.500	0.316	0.368
CenterFusion [7]	R&C	0.449	0.326	0.631	0.261	0.516	0.614	0.115
CenterNet [45]	C	0.400	0.338	0.658	0.255	0.629	1.629	0.142
FCOS3D [1]	C	0.428	0.358	0.690	0.249	0.452	1.434	0.124
PGD [20]	C	0.448	0.386	0.626	0.245	0.451	1.509	0.127
BEVDet [14]	C	0.482	0.422	0.529	0.236	0.395	0.979	0.152
PETR [13]	C	0.483	0.445	0.627	0.249	0.449	0.927	0.141
DD3D \dagger [42]	C	0.477	0.418	0.572	0.249	0.368	1.014	0.124
DETR3D \dagger [25]	C	0.479	0.386	0.626	0.245	0.394	0.845	0.133
BEVDet \dagger [14]	C	0.488	0.424	0.524	0.242	0.373	0.950	0.148
PETR \dagger [13]	C	0.504	0.441	0.593	0.249	0.383	0.808	0.132
Ours\dagger	R&C	0.517	0.453	0.569	0.246	0.379	0.781	0.128
Improvement	vs. R&C	+6.8%	+12.7%	+6.2%	+1.5%	+13.7%	-16.7%	-1.3%
Improvement	vs. C	+1.3%	+1.2%	+2.4%	+0.3%	+0.4%	+2.7%	+0.4%

TABLE II

COMPARISON OF SINGLE-FRAME STATE-OF-THE-ART WORKS ON THE nuSCENES VAL SET WITH DIFFERENT BACKBONES AND MODALS. † DENOTES ADOPTING THE DD3D [42] PRE-TRAINED V2-99 [43] BACKBONE.

Method	Backbone	Modal	NDS \uparrow	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
CenterNet [45]	DLA	C	0.328	0.306	0.716	0.264	0.609	1.426	0.658
CenterFusion [7]	DLA	R&C	0.453	0.332	0.649	0.263	0.535	0.540	0.142
Ours	DLA	R&C	0.471	0.419	0.664	0.258	0.416	0.789	0.193
		Improvement	+1.8%	+8.7%	-1.5%	+0.5%	+11.9%	-24.9%	-5.1%
BEVDet [14]	Res-50	C	0.381	0.304	0.719	0.272	0.555	0.903	0.257
PETR [13]	Res-50	C	0.381	0.313	0.768	0.278	0.564	0.923	0.225
Ours-R50	Res-50	R&C	0.402	0.330	0.712	0.255	0.552	0.898	0.214
		Improvement	+2.1%	+1.7%	+5.6%	+2.3%	+1.2%	+2.5%	+1.1%
FCOS3D [1]	Res-101	C	0.415	0.343	0.725	0.263	0.422	1.292	0.153
PGD [20]	Res-101	C	0.428	0.369	0.683	0.260	0.439	1.268	0.185
DETR3D [25]	Res-101	C	0.434	0.349	0.716	0.268	0.379	0.842	0.200
PETR [13]	Res-101	C	0.442	0.370	0.711	0.267	0.383	0.865	0.201
Ours	Res-101	R&C	0.455	0.380	0.675	0.258	0.372	0.833	0.196
		Improvement	+1.3%	+1.0%	+3.6%	+0.9%	+1.1%	+3.2%	+0.5%
PETR \dagger [13]	V2-99	C	0.455	0.403	0.736	0.271	0.432	0.825	0.204
Ours\dagger	V2-99	R&C	0.475	0.421	0.686	0.269	0.410	0.793	0.203
		Improvement	+2.0%	+1.8%	+5.0%	+0.2%	+2.2%	+3.2%	+0.1%

$[-10m, 10m]$ for Z axis. We set AdamW optimizer with the weight decay of 0.01 and initialize the learning rate with 2.0×10^{-4} . We select the 1600×640 resolution for nuScenes test set and different resolutions for nuScenes validation set. We randomly rotate the inputs in the range $[-22.5^\circ, 22.5^\circ]$. We train all experiments for 24 epochs with a total batch size of 8 on 8 RTX A6000 GPUs.

C. Comparisons with State-of-the-art Methods

Table I shows the single-frame 3D object detection performance comparison on nuScenes **test** set. Our method achieves the best performance on both NDS and mAP, compared with two Lidar-based methods InfoFocus [44], PointPillars [9], the radar-camera fusion approach CenterFusion [7], the camera-based monocular methods CenterNet [45], FCOS3D [1], PGD [20], DD3D [42] and the camera-based multi-view methods DETR3D [25], BEVDet [14] and PETR [13]. For the same radar-camera fusion, our approach exceeds the method CenterFusion [7] by 12.7% in mAP and 6.8% in NDS, respectively. The CenterFusion [7] performs the two-stage fusion approach which proposes two heads to regress velocity and attributes, so their mAVE and mAAE achieve better results. Then, compared with the latest camera

method PETR [13], our approach achieves state-of-the-art performance: 1.2% mAP and 1.3% NDS.

Furthermore, we compare the results with state-of-the-art methods on nuScenes **val** set, as shown in Table II. Following [13], we compare our approach to previous methods with different backbones. It shows that our approach also achieves the best performance on both NDS and mAP. Our approach surpasses the radar-camera fusion method CenterFusion [7] by 1.8 % NDS and 8.7 % mAP. Our method also outperforms the camera-based methods by evenly about 3% in mAP and about 2% to 3% in NDS. The excellent performance on the nuScenes test and validation set strongly validates our approach’s effectiveness.

D. Ablation Study

We perform ablation studies to validate the proposed modules, where we conduct all experiments using V2-99 [43] on nuScenes val set [46].

1) *Importance of proposed components.*: Table III shows the ablation studies about our two modules: semantic-aligned radar encoder (SARE) and radar-guided fusion transformer (RGFT), which perform improvement on the detection performance. Specifically, when inserting the SARE module

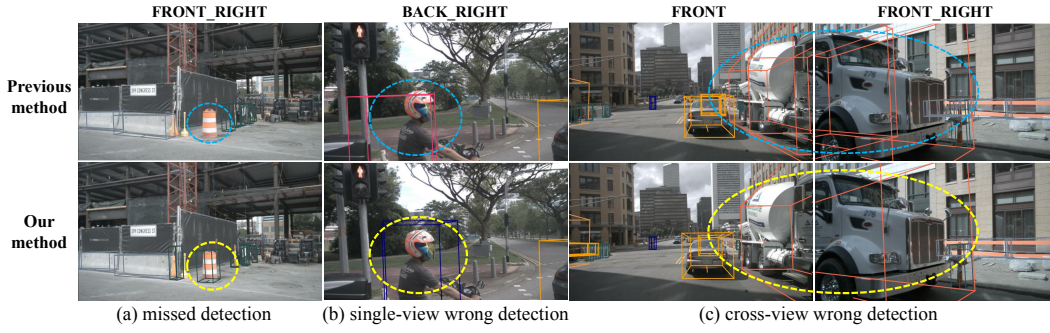


Fig. 6. The surround-view detection results comparison between our approach and the previous approach [13]. We mark the comparison areas with yellow circle for our approach and blue circle for [13], respectively. Our approach achieves correct object detection for different views, where our sufficient radar-camera interaction between semantic-aligned radar features and visual features performs more helpful clues for 3D detection.

TABLE III

ABLATION STUDY ON VAL SET FOR PROPOSED COMPONENTS. ‘SARE’ DENOTES THE SEMANTIC-ALIGNED RADAR ENCODER, ‘RGFT’ DENOTES THE RADAR-GUIDED FUSION TRANSFORMER.

	SARE	RGFT	NDS \uparrow	mAP \uparrow
(a)	-	-	0.457	0.406
(b)	✓	-	0.467	0.415
(c)	-	✓	0.465	0.414
(d)	✓	✓	0.475	0.421

TABLE IV

ABLATION STUDY ON VAL SET FOR SEMANTIC-ALIGNED RADAR ENCODER (SARE). ‘SI’ DENOTES THE SEMANTIC INDICATOR. ‘IGT’ DENOTES THE IMAGE-GUIDED RADAR TRANSFORMER.

	SI	IGT	NDS \uparrow	mAP \uparrow
(a)	-	-	0.465	0.414
(b)	✓	-	0.470	0.419
(c)	-	✓	0.468	0.417
(d)	✓	✓	0.475	0.421

(setting (b)), the semantic alignment and self-attention mechanism help achieve the semantic alignment for original sparse radar features to enhance the radar features’ representations. For RGFT (setting (c)), the fusion transformer brings global-scope adequate interaction, which reveals beneficial for detection. Finally, our method achieves the highest accuracy when applying both two modules.

2) *Ablation study for semantic-aligned radar encoder (SARE)*: We discuss the ablation effects within the proposed semantic-aligned radar encoder (SARE), as shown in Table IV. We remove all the SARE’s settings as the baseline of Table IV (setting (a)), which applies the radar inputs directly to the radar feature extractor with suboptimal grades. The setting (b) validates that the semantic indicator succeeds in employing the foreground’s semantics and relative location to obtain the alignment. And setting (c) reveals the image-guided radar transformer (IGRT), which promotes the long-range relationships for coarse radar features, performs a superior performance. All settings of SARE (setting (d)) achieve the most enhancement.

3) *Ablation study for radar-guided fusion transformer (RGFT)*: We investigate different settings within the radar-guided fusion transformer (RGFT), as shown in Table V. We adopt the setting (a) without RGFT as the baseline of Table V. Then setting (b) only utilizes the radar features as the RGFT’s query, key and value, actually changing the RGFT

TABLE V

ABLATION STUDY ON VAL SET FOR RADAR-GUIDED FUSION TRANSFORMER (RGFT). ‘w’ DENOTES ‘WITH’ AND ‘w/o’ DENOTES ‘WITHOUT’. ‘Q’, ‘K’, ‘V’ DENOTES QUERY, KEY, VALUE. ‘IMG.’ DENOTES IMAGE. ‘CONCAT.’ DENOTES ‘CONCATENATION’.

	Ablation	NDS \uparrow	mAP \uparrow
(a)	w/o RGFT	0.467	0.415
(b)	radar w/o img.	0.468	0.416
(c)	radar as Q w img. as K,V	N/A	N/A
(d)	radar concat. img. as Q w img. as K,V	0.475	0.421

to the self-attention, which improves limited performance, for we have conducted the self-attention for radar features in the above SARE. So we bring the image features as key and value of RGFT (setting (c)), but the sparse radar features influence the RGFT’s convergence and result in unstable training with exploding gradient. Finally, we combine the radar and image features as queries and adopt the image features as key and value (setting (d)), to strengthen the two modal’s interaction from the global scope to reach the best performance.

E. Visualization

We demonstrate the surround-view 3D detection visualization between our method and the previous approach [13] in Figure 6. Our approach achieves correct detection for the further, truncated, cross-view objects and releases the repeated prediction boxes. The semantic-aligned radar features offer more helpful clues and supplement the single-modal visual features to conduct more robust 3D detection.

V. CONCLUSION

In this paper, we provide a novel **Multi-View radar-camera Fusion** method **MVFusion** for 3D object detection, which achieve the semantic-aligned radar features and robust cross-modal information interaction. Specifically, we propose the semantic-aligned radar encoder (SARE) to extract image-guided radar features. After radar feature extraction, we propose the radar-guided fusion transformer (RGFT) to integrate the enhanced radar features with high-level image features. Extensive experiments on the nuScenes dataset validate that our model achieves single-frame radar-camera fusion state-of-the-art performance. In the future, we will aggregate the spatio-temporal information from multi-view cameras, to further promote the radar-camera fusion.

REFERENCES

- [1] T. Wang, X. Zhu, J. Pang, and D. Lin, "FCOS3D: Fully convolutional one-stage monocular 3D object detection," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2021, pp. 913–922.
- [2] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "BEVFusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," *arXiv preprint arXiv:2205.13542*, 2022.
- [3] M. Liang, B. Yang, S. Wang, and R. Urtasun, "Deep continuous fusion for multi-sensor 3D object detection," in *European Conf. on Computer Vision (ECCV)*, 2018, pp. 641–656.
- [4] Y. Wang, Z. Jiang, X. Gao, J.-N. Hwang, G. Xing, and H. Liu, "RODNet: Radar object detection using cross-modal supervision," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2021, pp. 504–513.
- [5] K. Qian, S. Zhu, X. Zhang, and L. E. Li, "Robust multimodal vehicle detection in foggy weather using complementary lidar and radar signals," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 444–453.
- [6] Y. Li, J. Deng, Y. Zhang, J. Ji, H. Li, and Y. Zhang, "EZFusion: A close look at the integration of lidar, millimeter-wave radar, and camera for accurate 3d object detection and tracking," *IEEE Robotics and Automation Letters (RAL)*, 2022.
- [7] R. Nabati and H. r. Qi, "CenterFusion: Center-based radar and camera fusion for 3D object detection," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2021, pp. 1527–1536.
- [8] F. Nobis, M. Geisslinger, M. Weber, J. Betz, and M. Lienkamp, "A deep learning-based radar and camera sensor fusion architecture for object detection," in *2019 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*. IEEE, 2019, pp. 1–7.
- [9] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "PointPillars: Fast encoders for object detection from point clouds," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2019, pp. 12 697–12 705.
- [10] X. Dong, B. Zhuang, Y. Mao, and L. Liu, "Radar camera fusion via representation learning in autonomous driving," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 1672–1681.
- [11] R. Nabati and H. Qi, "Radar-camera sensor fusion for joint object detection and distance estimation in autonomous vehicles," *arXiv preprint arXiv:2009.08428*, 2020.
- [12] Y. Kim, J. W. Choi, and D. Kum, "GRIF Net: Gated region of interest fusion network for robust 3d object detection from radar point cloud and monocular image," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 857–10 864.
- [13] Y. Liu, T. Wang, X. Zhang, and J. Sun, "PETR: Position embedding transformation for multi-view 3d object detection," in *European Conf. on Computer Vision (ECCV)*, 2022.
- [14] J. Huang, G. Huang, Z. Zhu, and D. Du, "BEVDet: High-performance multi-camera 3D object detection in bird-eye-view," *arXiv preprint arXiv:2112.11790*, 2021.
- [15] O. Schumann, M. Hahn, N. Scheiner, F. Weishaupt, J. F. Tilly, J. Dickmann, and C. Wöhler, "RadarScenes: A real-world radar point cloud data set for automotive applications," in *International Conference on Information Fusion (FUSION)*. IEEE, 2021, pp. 1–8.
- [16] P. Li, H. Zhao, P. Liu, and F. Cao, "RTM3D: Real-time monocular 3D detection from object keypoints for autonomous driving," in *European Conf. on Computer Vision (ECCV)*. Springer, 2020, pp. 644–660.
- [17] P. Li and H. Zhao, "Monocular 3D detection with geometric constraint embedding and semi-supervised training," *IEEE Robotics and Automation Letters (RAL)*, vol. 6, no. 3, pp. 5565–5572, 2021.
- [18] Y. Liu, Y. Yixuan, and M. Liu, "Ground-aware monocular 3D object detection for autonomous driving," *IEEE Robotics and Automation Letters (RAL)*, vol. 6, no. 2, pp. 919–926, 2021.
- [19] C. Reading, A. Harakeh, J. Chae, and S. L. Waslander, "Categorical depth distribution network for monocular 3D object detection," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2021, pp. 8555–8564.
- [20] T. Wang, Z. Xinge, J. Pang, and D. Lin, "Probabilistic and geometric depth: Detecting objects in perspective," in *Conf. on Robot Learning (CoRL)*. PMLR, 2022, pp. 1475–1485.
- [21] Y. Zhou and O. Tuzel, "VoxelNet: End-to-end learning for point cloud based 3D object detection," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4490–4499.
- [22] Y. Yan, Y. Mao, and B. Li, "SECOND: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, 2018.
- [23] S. Shi, X. Wang, and H. Li, "PointRCNN: 3D object proposal generation and detection from point cloud," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 770–779.
- [24] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3D," in *European Conf. on Computer Vision (ECCV)*. Springer, 2020, pp. 194–210.
- [25] Y. Wang, V. C. Guizilini, T. Zhang, Y. Wang, H. Zhao, and J. Solomon, "DETR3D: 3d object detection from multi-view images via 3d-to-2d queries," in *Conf. on Robot Learning (CoRL)*. PMLR, 2022, pp. 180–191.
- [26] B. Pan, J. Sun, H. Y. T. Leung, A. Andonian, and B. Zhou, "Cross-view semantic segmentation for sensing surroundings," *IEEE Robotics and Automation Letters (RAL)*, vol. 5, no. 3, pp. 4867–4873, 2020.
- [27] Z. Liu, Z. Wu, and R. Tóth, "SMOKE: Single stage monocular 3D object detection via keypoint estimation," in *IEEE Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 996–997.
- [28] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Q. Yu, and J. Dai, "BEVFormer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers," in *European Conf. on Computer Vision (ECCV)*, 2022.
- [29] Y. Long, D. Morris, X. Liu, M. Castro, P. Chakravarty, and P. Narayanan, "Full-velocity radar returns by radar-camera fusion," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2021, pp. 16 198–16 207.
- [30] J.-T. Lin, D. Dai, and L. Van Gool, "Depth estimation from monocular images and sparse radar data," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 233–10 240.
- [31] B. Yang, R. Guo, M. Liang, S. Casas, and R. Urtasun, "RadarNet: Exploiting radar for robust perception of dynamic objects," in *European Conf. on Computer Vision (ECCV)*. Springer, 2020, pp. 496–512.
- [32] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3D object detection network for autonomous driving," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1907–1915.
- [33] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. L. Waslander, "Joint 3D proposal generation and object detection from view aggregation," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–8.
- [34] S. Vora, A. H. Lang, B. Helou, and O. Beijbom, "PointPainting: Sequential fusion for 3D object detection," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 4604–4612.
- [35] Y. Li, A. W. Yu, T. Meng, B. Caine, J. Ngiam, D. Peng, J. Shen, Y. Lu, D. Zhou, Q. V. Le, et al., "DeepFusion: Lidar-camera deep fusion for multi-modal 3D object detection," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 17 182–17 191.
- [36] X. Bai, Z. Hu, X. Zhu, Q. Huang, Y. Chen, H. Fu, and C.-L. Tai, "TransFusion: Robust Lidar-camera fusion for 3D object detection with transformers," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 1090–1099.
- [37] C. Wang, C. Ma, M. Zhu, and X. Yang, "Pointaugmenting: Cross-modal augmentation for 3D object detection," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 11 794–11 803.
- [38] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [39] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., "AN IMAGE IS WORTH 16x16 WORDS: Transformers for image recognition at scale," in *Int. Conf. on Learning Representations (ICLR)*, 2021.
- [40] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [41] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2017, pp. 2980–2988.
- [42] D. Park, R. Ambrus, V. Guizilini, J. Li, and A. Gaidon, "Is pseudo-lidar needed for monocular 3d object detection?" in *IEEE Int. Conf. on Computer Vision (ICCV)*. IEEE, 2021, pp. 3142–3152.
- [43] Y. Lee and J. Park, "CenterMask: Real-time anchor-free instance segmentation," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 13 906–13 915.
- [44] J. Wang, S. Lan, M. Gao, and L. S. Davis, "InfoFocus: 3d object detection for autonomous driving with dynamic information modeling,"

- in *European Conf. on Computer Vision (ECCV)*. Springer, 2020, pp. 405–420.
- [45] X. Zhou, D. Wang, and P. Krähenbühl, “Objects as points,” *arXiv preprint arXiv:1904.07850*, 2019.
- [46] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuScenes: A multimodal dataset for autonomous driving,” in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2020, pp. 11 621–11 631.
- [47] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2016, pp. 770–778.
- [48] Y. Tang, S. Dorn, and C. Savani, “Center3D: Center-based monocular 3D object detection with joint depth understanding,” in *German Conf. on Pattern Recognition (GCPR)*. Springer, 2020, pp. 289–302.