

RC-BEV Fusion: A Plug-In Module for Radar-Camera Bird’s Eye View Feature Fusion

Lukas Stacker^{1,2}[0000–0001–7537–606X], Shashank Mishra^{2,3}[0000–0003–2882–2789],
Philipp Heidenreich¹[0000–0002–5906–7125], Jason Rambach³[0000–0001–8122–6789],
and Didier Stricker^{2,3}[0000–0002–5708–6023]

¹ Stellantis, Opel Automobile GmbH, Russelsheim am Main, Germany

² Rheinland-Pfalzische Technische Universitat Kaiserslautern-Landau,
Kaiserslautern, Germany

³ German Research Center for Artificial Intelligence, Kaiserslautern, Germany

Abstract. Radars and cameras belong to the most frequently used sensors for advanced driver assistance systems and automated driving research. However, there has been surprisingly little research on radar-camera fusion with neural networks. One of the reasons is a lack of large-scale automotive datasets with radar and unmasked camera data, with the exception of the nuScenes dataset. Another reason is the difficulty of effectively fusing the sparse radar point cloud on the bird’s eye view (BEV) plane with the dense images on the perspective plane. The recent trend of camera-based 3D object detection using BEV features has enabled a new type of fusion, which is better suited for radars. In this work, we present RC-BEV Fusion, a modular radar-camera fusion network on the BEV plane. We propose two novel radar encoder branches, and show that they can be incorporated into several state-of-the-art camera-based architectures. We show significant performance gains of up to 28% increase in the nuScenes detection score, which is an important step in radar-camera fusion research. Without tuning our model for the nuScenes benchmark, we achieve the best result among all published methods in the radar-camera fusion category.

1 Introduction

The development of advanced driver assistance systems (ADAS) and automated driving functions has made remarkable progress in recent years, resulting in increased safety and convenience for drivers. A robust environment perception is the key requirement for these systems, which rely on sensors such as radar, camera, or LiDAR to detect surrounding objects. Each sensor has its own advantages and disadvantages that must be considered when designing a perception system. An extensive review on multi-modular automotive object detection is presented in [4].

Radar sensors are advantageous in that they are less affected by adverse environmental conditions such as rain, fog, or darkness, and they have a longer detection range when compared to cameras and LiDARs. However, they are

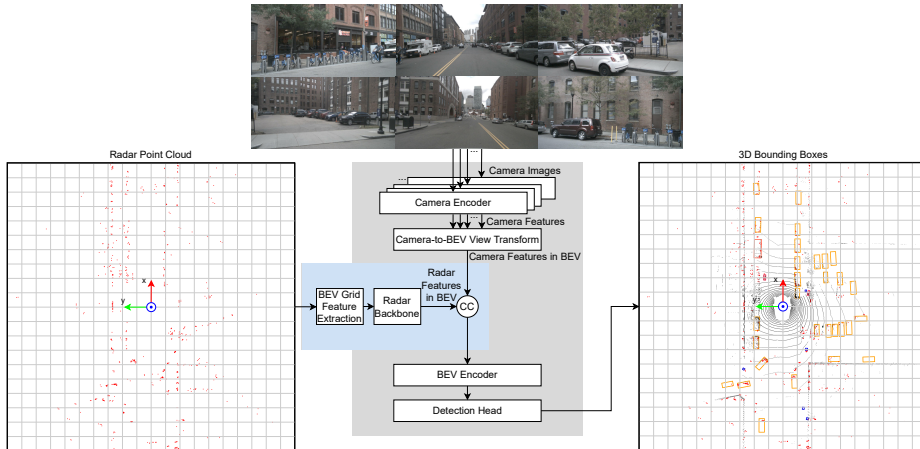


Fig. 1: Overview of RC-BEVFusion network architecture. The block marked in grey is inherited from an exchangeable camera-only baseline, while the block marked in blue shows our proposed radar-camera fusion plug-in module.

limited in their ability to provide detailed information about the shape and texture of objects [58]. Cameras, on the other hand, provide rich visual information and can recognize objects based on their appearance, but their performance can be affected by changes in lighting conditions and inaccurate depth estimation [11, 24]. LiDARs provide detailed 3D information and are less affected by lighting conditions, but they can be expensive and have limited range [58].

Sensor fusion has the potential to overcome these limitations of individual sensors. In particular, the combination of radar and camera sensors arguably offers the most complementary features. The main challenge is how to associate radar and camera features given that conventional radars provide data on the bird’s eye view (BEV) plane, whereas cameras provide data on the image plane. Projecting radar points to the image discards too much geometric information, whereas projecting camera features to the sparse radar points discards too much semantic information [22].

Recent advancements in camera-only networks using view transformers [10, 36] have enabled a new type of fusion on the BEV plane, which is well suited for radar data. In this paper, we propose RC-BEVFusion, a novel radar-camera fusion architecture on the BEV plane inspired by [22] and illustrated in Figure 1. In contrast to previous radar-camera fusion techniques [14, 31, 35], our architecture allows radar and camera features to equally contribute to the final detections, enabling the network to detect obstacles that may be missed by one of the modalities. It is a flexible architecture that inherits several elements from an exchangeable camera-only baseline: a camera encoder, a camera-to-BEV view transformer, a BEV encoder and a detection head. On top of these modules, we propose two radar encoder branches: RadarGridMap and BEVFeatureNet. Our

results show that they can be used as a plug-in module in various camera-based architectures and significantly enhance their performance.

To train and evaluate our network, we need a large-scale automotive dataset with radar point clouds, unmasked camera images with a lot of variety in the scenes and 3D object annotations. A recent overview on radar datasets is given in [58]. First, there are datasets with conventional 2+1D radar sensors that provide a list of detections with measurements of the range, range rate, azimuth angle, and radar cross section (RCS). Out of these, the nuScenes dataset [1] is the only one that fulfils our requirements. Second, there are recent datasets with high-performance 3+1D radar sensors that provide denser point clouds and additionally measure the elevation angle. From this group, the Astyx [27] dataset is too small, while the View-of-Delft [34] dataset is medium-sized but has limited visual variety in the images due to the high annotation frequency. The recently presented TJ4DRadSet [55] may be a future option but the data is not fully released yet. In this work, we therefore choose to conduct our experiments with the nuScenes dataset [1].

The remaining paper is organized as follows. Section 2 provides an overview of related work on automotive object detection with camera-only, radar-only and radar-camera fusion. Section 3 describes the proposed radar-camera fusion architecture on the BEV plane. Section 4 presents extensive experimental results on the nuScenes dataset and demonstrates the effectiveness of the proposed architecture. Finally, Section 5 concludes the paper and outlines future work.

2 Related work

The task of 3D object detection is mostly conducted with the help of cameras, LiDARs and, less frequently, radars. In the following, we give an overview of recent advances on image-only and radar-only object detection, as well as LiDAR-camera and radar-camera sensor fusion.

2.1 Image-only object detection

Image-based 3D object detection is a difficult task in the field of computer vision because it involves identifying and localizing objects in 3D space using only a single camera as a sensor. This is in contrast to LiDAR and radar systems, which provide depth measurements and can more accurately determine the 3D location of objects.

Early approaches to image-based 3D object detection focused on using known geometric information to estimate 3D bounding boxes from 2D detections [28]. More recent techniques have extended existing 2D object detection models with additional detection heads specifically designed for 3D object detection [43, 48, 57]. Some approaches have also used predicted depth maps as auxiliary features [26] or to create a pseudo-LiDAR point cloud, which is then processed using LiDAR-based object detectors [49].

The most recent research in this area has mainly followed two directions. The first is the use of transformer-based techniques, which leverage the ability of transformer models to process sequences of data and perform self-attention to learn more complex relationships between features [20, 50]. The second direction is the development of BEV-based object detectors. To this end, the features need to be transformed from the image plane to the BEV plane. A pioneering work uses orthographic feature transform [39], where a voxel grid is projected to the image to extract features. To reduce memory consumption, the voxel grid is then collapsed along the vertical axis to create BEV features. More recent methods are based on the Lift-Splat-Shoot (LSS) view transformer [36], which lifts image features into a 3D pseudo point cloud via dense depth prediction, before again collapsing the vertical dimension to create BEV features. The idea is first incorporated into a 3D object detection network in [10], and later refined with LiDAR-based depth supervision [18] and temporal stereo [17]. In [56], a more efficient view transformer in terms of memory and computations is introduced by formulating LSS into matrix operations with decomposed ring and ray matrices, and compressing image features and the estimated depth map along the vertical dimension. In this work, we build upon these BEV-based object detectors and integrate our radar-camera fusion as a plug-in module.

2.2 Radar-only object detection

Automotive radars are a common sensor used in autonomous vehicles for detecting objects in the environment. These radars typically provide a preprocessed list of detections, but the sparsity and lack of semantic information make it difficult to use the data for stand-alone 3D object detection. As a result, much of the research in this area has focused on either semantic segmentation of radar point clouds [42] or experimental setups using the raw radar cube [25, 33].

Recently, there have been some point cloud-based techniques for object detection as well. These can be broadly divided into convolutional and graph-based approaches. Some convolutional approaches assign each point in the point cloud to a cell in a BEV grid and include feature layers such as the maximum RCS and Doppler values [5, 40]. Since the BEV grid is similar in structure to an image, traditional convolutional networks can be applied for object detection. Other convolutional approaches use variants of PointPillars [15] to create a pillar grid automatically from the point cloud [34, 40]. In contrast, graph neural network based approaches perform object detection directly on the radar point cloud [3, 41]. Recent work combines both approaches by first extracting features with a graph neural network and then mapping the features to a BEV grid for further processing [45]. In this work, we examine radar feature encoders inspired by these ideas.

2.3 Sensor fusion object detection

Sensor fusion aims at leveraging the strengths of a diverse sensor combination. Most sensor fusion research focuses on LiDAR-camera fusion, as LiDAR provides

accurate 3D information and cameras provide high semantic value, while sharing the same optical propagation principles. There are several techniques for fusing LiDAR and camera data. Some approaches are based on projecting 2D detections from the camera into a frustum and matching them with LiDAR points to refine the 3D detections [37, 51]. Other techniques augment the LiDAR points with semantics from the image and use LiDAR-based object detectors to perform detection [7, 46, 47]. The most recent approaches to LiDAR-camera fusion involve extracting BEV features from both the LiDAR and camera data and fusing them on the BEV plane before applying a joint BEV encoder to perform object detection [6, 19, 22].

When compared to LiDAR, automotive radar uses a different wavelength and measurement principle. As a consequence, radar typically shows strong RCS fluctuations and reduced resolution in range and angle, resulting in less dense point clouds. Moreover, whereas modern radars also measure elevation, many conventional radars only provide detections on the BEV plane. These differences make the fusion especially challenging and prevent the simple replacement of LiDAR with radar processing. Early research commonly projected the radar detections onto the image plane to associate the data. This approach can be used to find regions of interest in the image [29, 30] or to create additional image channels with radar data that can be used with image-based networks [2, 32, 44].

More recent methods have moved away from this 2D approach and instead focus on fusing based on 3D information. One approach is to refine image-based 3D detections with associated radar data [31]. This can be sub-optimal because it discards the possibility of radar-only detections. Another approach is to project 3D regions of interest (ROIs) to the image and BEV plane to extract features from each sensor [13]. Finally, cross-attention has been used to align and fuse the features in 3D [12, 35]. In this work, we propose a novel architecture to fuse radar and camera data on the BEV plane.

3 RC-BEV Fusion

In this section, we describe our proposed model architectures. We start by giving an overview of the general fusion architecture, before providing more detail on the proposed radar encoders, the camera-only networks we use as baselines and the loss function.

3.1 Overview

In this work, we introduce a novel radar branch and use it as a plug-in module on different camera-based 3D object detection networks to improve their performance. The prerequisite for our proposed radar-camera fusion is that the camera-only network uses BEV features as an intermediate representation. The general architecture is shown in Figure 1. The block marked in grey is inherited from an exchangeable camera-only baseline, while the block marked in blue shows our proposed radar plug-in module.




Fig. 2: BEVFeatureNet radar encoder. The radar detections are mapped to BEV grid cells for augmentation and restructuring into a dense tensor. After applying a simplified PointNet, the encoded features are remapped to the BEV grid.

First, BEV features are extracted from the images and the radar point cloud separately. To this end, a backbone network is used to extract features from each image before they are transformed into joint BEV features with a view transformer. We set up our radar encoder so that it creates BEV features in the same shape and geometric orientation as the camera BEV features. The features are then fused by concatenation followed by a 1×1 convolutional layer that reduces the embedding dimension to the original dimension of the camera BEV features. We can then use the same BEV encoder and 3D detection heads that are used in the respective camera-only network, introducing little overhead with our fusion.

3.2 Radar encoders

We propose two radar encoders, RadarGridMap and BEVFeatureNet, which we explain in the following. They each consist of two stages: First, we create a regular, structured BEV grid from the sparse radar point cloud. Then, we apply a convolutional backbone that further encodes the BEV features.

RadarGridMap Inspired by [5], we design a hand-crafted radar BEV grid. We map each detection to a cell on the grid and fill the cell with four channels: the number of detections per cell, the maximum RCS value, and the minimum and maximum signed compensated Doppler values.

After the grid mapping, we use a small generalized ResNet [8] as our radar backbone. We use 16 layers grouped into residual blocks with BatchNorm and ReLU. We use two downsampling stages that double the channels but reduce the resolution of the BEV grid. We design the size of the BEV grid so that the output of the radar backbone has the same shape as the camera BEV features.

BEVFeatureNet The BEVFeatureNet illustrated in Figure 2 is inspired by the pillar feature encoding of PointPillars [15], but adapted for radar data. First, we map each point of the radar point cloud to a cell in a predefined BEV grid. In the original pillar feature encoding, each point in the LiDAR point cloud has

Table 1: Camera-only network configurations. For FPN-LSS, SECOND-FPN and Generalized ResNet, we indicate the number of output channels of the module. For LSS and MatrixVT, we indicate the number of channels and the resolution of the output BEV feature in meters. MF denotes multi-frame temporal fusion.

Model	Camera Encoder	View Transf.	BEV Encoder	Head	Input Res.	MF
BEVDet [10]	Swin-Tiny	LSS	Gen. ResNet-512	CenterPoint	704x256	
	FPN-LSS-256	80-0.4-0.4	FPN-LSS-256			
BEVDepth [18]	ResNet-50	LSS	ResNet-18	CenterPoint	704x256	✓
	SECOND-FPN-128	80-0.8-0.8	SECOND-FPN-64			
BEVStereo [17]	ResNet-50	LSS	ResNet-18	CenterPoint	704x256	✓
	SECOND-FPN-128	80-0.8-0.8	SECOND-FPN-64			
MatrixVT [56]	ResNet-50	MatrixVT	ResNet-18	CenterPoint	704x256	
	SECOND-FPN-128	80-0.8-0.8	SECOND-FPN-64			

coordinates x , y , and z , and reflectivity r . Further, the point is augmented with the distances to the arithmetic mean of all points in its pillar, x_c , y_c , and z_c and the offsets to the pillar center x_p and y_p . The 2+1D radar point cloud in nuScenes [1] does not have a z -coordinate or a reflectivity r , but instead it has a radial velocity v_d , measured via the Doppler effect, and an RCS. We therefore discard the z -axis and its augmented feature, replace the reflectivity with the RCS, and include the radial velocity values. We further aggregate multiple radar sweeps and append the timestamp difference to the latest sweep t_s to each point. Thus, we obtain the 9-dimensional feature set:

$$\vec{F} = [x, y, \text{RCS}, v_d, t_s, x_c, y_c, x_p, y_p] \quad (1)$$

As in [15], we then use a set of non-empty BEV grid cells B with a fixed number of points per cell N_p to create a dense tensor of size (F, B, N_p) . If the number of non-empty BEV grid cells or the number of points per cell is lower or higher than the fixed number, we apply zero-padding or random sampling, respectively. For each point, we then apply a simplified PointNet [38] with a 1×1 convolutional layer followed by BatchNorm and ReLU resulting in a tensor of shape (C, B, N_p) , before a max operation over the points per cell reduces the dimension to (C, B) . We then map the C -dimensional features back to their position on the BEV grid. Finally, the same convolutional backbone as for RadarGridMap is applied.

3.3 Camera-only baselines

We selected various camera-only baselines to showcase the plug-in character of our radar fusion module. In this section, we list details for the camera baselines we examined. A compact overview of the modules is given in Table 1.

BEVDet BEVDet [10] is the first network that uses BEV-based features for object detection on the nuScenes dataset. First, high-level features from each of the N_i input images of shape (H_i, W_i) are extracted separately. To this end, a SwinTransformer-Tiny [21] backbone network outputs multi-scale feature maps, which are then processed using the feature pyramid network from [36], FPN-LSS,

which upsamples the low resolution feature maps to match the high resolution map, concatenates them and runs them through a ResNet [8] block. This leads to a feature map of shape $(N_i, H_i/8, W_i/8, C)$ with C feature channels. Then, a 1×1 convolution followed by a Softmax is used to predict a depth classification into D pre-defined depth bins. An outer product across the feature and depth classification channels creates a very large tensor of shape $(N_i, H_i/8, W_i/8, D, C)$. Using the intrinsic and extrinsic calibration matrices of each camera, this tensor can be unprojected into a pseudo-pointcloud. The vertical dimension of this pseudo-pointcloud is then collapsed by summing up the features from all points that fall into the same cell of a pre-defined BEV grid with shape (H_{bev}, W_{bev}) . We follow the implementation of [22], which provides computationally efficient BEV pooling and uses a heavier view transformer to enable more accurate depth estimation, which is important to associate the features from radar and camera. To further encode the BEV features, the generalized ResNet [8] structure from [36] followed again by FPN-LSS is used.

BEVDepth BEVDepth [18] uses a similar structure to BEVDet [10] but aims to achieve a more accurate depth estimation. To this end, the single convolutional layer for depth estimation in BEVDet [10] is replaced with a camera-aware DepthNet module, that concatenates and flattens the camera’s intrinsic and extrinsic calibration parameters and uses an MLP to rescale them to match the dimension of the image features C . This calibration vector is used to re-weight the image features using a Squeeze-and-Excitation module [9]. During training, BEVDepth [18] further uses the depth value from projected LiDAR points on the image to directly supervise the depth estimation via binary cross-entropy loss. We use the released configuration, which encodes the images with a ResNet-50 [8] backbone followed by the feature pyramid net from [52], SECOND-FPN, which concatenates upsampled multi-scale feature maps. The BEV encoder uses a ResNet-18 again followed by SECOND-FPN. The view transformer has a lower resolution than the one we use for BEVDet [10]. It also uses a multi-frame fusion with one previous keyframe, thus two images from each camera taken 500 ms apart are used to create the BEV features, resulting in more accurate velocity estimation.

BEVStereo BEVStereo [17] builds upon BEVDepth [18] and uses the same configuration for most modules. In addition to the monocular depth estimation with DepthNet, it introduces a temporal stereo depth estimation module, which is based on multi-view-stereo [53]. To this end, for each pixel in the current image feature map, several depth candidates are predicted and used to retrieve corresponding features from the previous image features using a homography warping. The confidence of each candidate is evaluated based on the current and previous feature’s similarity and used to iteratively optimize the depth candidates. After three iterations, the depth candidates are used to construct the stereo depth. Since the stereo depth is not viable for pixels that do not have corresponding pixels in the previous image, a convolutional WeightNet module is

used to combine the monocular depth estimations from the current and previous image with the stereo depth estimation to produce the final depth estimates.

MatrixVT In MatrixVT [56], an alternative view transformer is proposed. The view transformation step of LSS [36] is first generalized into matrix operations, which leads to a very large and sparse feature transportation tensor of shape $(H_{bev}, W_{bev}, N_i, H_i/8, W_i/8, D)$, which transforms the image feature tensor with depth estimates to the BEV features. To combat this, the image feature and the dense depth prediction are compressed along the vertical dimension, resulting in prime feature and prime depth matrices. This is feasible due to the low response variance in the height dimension of the image. To further reduce its sparsity, the feature transportation tensor is orthogonally decomposed into a ring tensor, which encodes distance information, and a ray tensor, which encodes directional information. Using efficient mathematical operations, the resulting view transformer achieves lower computational cost and memory requirements. We use a configuration based on BEVDepth [18], which only replaces the view transformer with MatrixVT [56] and does not use multi-frame fusion.

3.4 Detection Head and Loss

All examined camera baselines use the CenterPoint [54] detection head, so we can apply the same loss in all architectures. For each class k , CenterPoint [54] predicts a BEV heatmap with peaks at object center locations. The heatmap is trained using a ground-truth heatmap y filled with Gaussian distributions around ground truth object centers. Given the heatmap score p_{kij} at position i, j in the BEV grid and the ground truth y_{kij} , we can compute the Gaussian focal loss [16] as:

$$L_{hm} = -\frac{1}{N_o} \sum_{k=1}^K \sum_{i=1}^{H_{bev}} \sum_{j=1}^{W_{bev}} \begin{cases} (1 - p_{kij})^\alpha \log(p_{kij}) & y_{kij} = 1 \\ (1 - y_{kij})^\beta (p_{kij})^\alpha \log(1 - p_{kij}) & \text{otherwise} \end{cases} \quad (2)$$

where N_o is the number of objects per image, K is the number of classes, H_{bev} and W_{bev} are the height and width of the BEV grid, and α and β are hyper-parameters. In addition, CenterPoint [54] has regression heads that output all parameters needed to decode 3D bounding boxes: a sub-pixel location refinement, a height above ground, the dimensions, the velocity, and the sine and cosine of the yaw rotation angle. The regression heads are trained with L1 loss.

4 Experimental Results

In this section, we present our experimental results. We first give some more information on the dataset, evaluation metrics and training settings. We then list quantitative results on the nuScenes validation set to ensure a fair comparison between our proposed fusion networks with their camera-only baselines. We also

show results for the nuScenes [1] test benchmark and provide an inference example for a qualitative comparison. Ablative studies, detailed class-wise results, rain and night scene evaluation and more qualitative examples are provided in the supplementary material.

4.1 Data and metrics

For our experiments, we need a large-scale, real-world dataset with unmasked camera images, series-production automotive radars and 3D bounding box annotations. To the best of our knowledge, the nuScenes dataset [1] is currently the only dataset that fulfils these requirements. It covers data from six cameras at 12 Hz, five conventional 2+1D radar sensors at 13 Hz and one LiDAR at 20 Hz, as well as 1.4 million 3D bounding boxes annotated at 2 Hz. We follow the official splits into 700 training, 150 validation and 150 test scenes and reduce the 27 annotated classes to the 10 classes evaluated on the benchmark. We also use the official metrics: the mean average precision (mAP), the true positive metrics covering mean errors for translation (mATE), scale (mASE), orientation (mAOE), velocity (mAVE) and nuScenes attribute (mAAE), as well as the condensed nuScenes detection score (NDS).

4.2 Quantitative Evaluation

Training details For our radar-camera fusion networks, we adopt the configurations from the camera-only baselines listed in Table 1 to allow for a fair comparison. In addition, we design our radar encoder branch so that the BEV features have the same shape and orientation as the camera BEV features. To increase the point cloud density while limiting the use of outdated data, we choose to aggregate five radar sweeps, which corresponds to 300 ms of past data. The aggregated radar point cloud is still sparse when compared with LiDAR data, so that we can reduce the number of non-empty grid cells and points per grid cell of the BEVFeatureNet drastically with respect to the pillar feature encoding in [15]. We empirically find that setting $B = 2000$, $N_p = 10$, and $C = 32$ is sufficient, which allows for little computational overhead. For the Gaussian focal loss, we follow [16] and set $\alpha = 2$ and $\beta = 4$. We train for 20 epochs with an AdamW [23] optimizer, a base learning rate of $2e-4$ and weight decay of $1e-2$.

NuScenes validation results We show the results for our radar-camera fusion networks w.r.t. their camera-only baselines on the nuScenes validation split in Table 2. First, we compare the two radar encoders with our model based on BEVDet. In both cases, the proposed fusion offers significant performance increases, with the mAP increasing up to 24% and the NDS 28%. The up to 23% reduced translation error shows how the direct depth measurements provided by the radar can lead to more precise location predictions. The most significant improvement of 55% is achieved for the velocity error, which is enabled by the direct velocity measurement of the radar. This effect also helps determining whether an

Table 2: Experimental results for our radar-camera fusion used in different architectures on the nuScenes val split. The inference latency T is measured on a Nvidia RTX 2080 Ti. *We use the implementation of BEVDet-Tiny with a heavier view transformer from [22]. †We list the results as reported by the authors.

	Cam. model	Radar model		mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow	T [ms]
[10]	BEVDet*	None		0.350	0.411	0.660	0.275	0.532	0.918	0.260	122
Ours	BEVDet*	RadarGridMap	Δ_r	0.429	0.525	0.523	0.272	0.507	0.412	0.183	132
				23%	28%	-21%	-1%	-5%	-55%	-30%	
Ours	BEVDet*	BEVFeatureNet	Δ_r	0.434	0.525	0.511	0.270	0.527	0.421	0.182	139
				24%	28%	-23%	-2%	-1%	-54%	-30%	
[18]†	BEVDepth	None		0.359	0.480	0.612	0.269	0.507	0.409	0.201	132
Ours	BEVDepth	BEVFeatureNet	Δ_r	0.405	0.521	0.542	0.274	0.512	0.309	0.181	146
				13%	9%	-11%	2%	1%	-24%	-10%	
[17]†	BEVStereo	None		0.372	0.500	0.598	0.270	0.438	0.367	0.190	308
Ours	BEVStereo	BEVFeatureNet	Δ_r	0.423	0.545	0.504	0.268	0.453	0.270	0.174	322
				14%	9%	-16%	-1%	3%	-26%	-8%	
[56]	MatrixVT	None		0.319	0.400	0.669	0.281	0.494	0.912	0.238	54
Ours	MatrixVT	BEVFeatureNet	Δ_r	0.386	0.495	0.549	0.275	0.539	0.423	0.193	64
				21%	24%	-18%	-2%	9%	-54%	-19%	

Table 3: Experimental results for published radar-camera fusion models on the nuScenes test benchmark.

Model	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
CenterFusion [31]	0.326	0.449	0.631	0.261	0.516	0.614	0.115
CRAFT [14]	0.411	0.523	0.467	0.268	0.456	0.519	0.114
TransCAR [35]	0.422	0.522	0.630	0.260	0.383	0.495	0.121
Ours (BEVDet)	0.476	0.567	0.444	0.244	0.461	0.439	0.128

object is currently moving or stopped, thus reducing the attribute error. The two metrics that remain relatively unaffected by the fusion are scale and orientation error. This is as expected since the sparse radar detections do not help to determine an object’s size or orientation. Overall, we observe similar results for the RadarGridMap and the BEVFeatureNet encoder. This demonstrates the effectiveness and modularity of the proposed BEV-based feature fusion. In general, we recommend using the BEVFeatureNet because it requires less hand-crafted input, is more scalable, and achieves slightly better results.

In the second part of Table 2, we use the BEVFeatureNet encoder as a plug-in branch in different camera-only baselines. We observe significant performance increase for all examined architectures, again confirming the modularity of the proposed architecture. There are two potential reasons for the difference in relative performance increase between the camera-only baselines. First, BEVDepth and BEVStereo use temporal fusion and therefore achieve better velocity prediction and overall scores, leading to smaller margins for the radar-camera fusion. Second, we use a BEVDet variant with a heavier view transformer especially designed for fusion on the BEV space. This modification may explain the relatively high performance gains.

Finally, we also measure the inference latency on an Nvidia RTX 2080 Ti GPU to demonstrate that our fusion approach introduces only small computational overhead due to the efficient radar encoder design.

4.3 Qualitative Evaluation

To get a better impression of the difference in object detection performance, we present an inference example for BEVDet [10] and RC-BEV Fusion based on BEVFeatureNet and BEVDet in Figure 3. We show the camera-only inference results and radar-camera fusion results in the left and right subfigure, respectively. We show the full surround view with all six cameras, the top row shows the front left, center and right camera, while the bottom row shows the back right, center and left camera. On the bottom, we show a BEV perspective with LiDAR points in black for reference and radar points in red for the fusion network only. In each perspective, we show the projected 3D bounding boxes predicted by the networks, where the color indicates the class: pedestrians are blue, cars are yellow and trucks are orange.

In the scene, we see a crowded intersection with lots of cars and pedestrians. At first, the visual impression when looking at the camera images is that most objects are well detected. However, comparing the dashed red ellipses in the BEV perspective on the right, we can see that the radar-camera fusion enables much more accurate detection of the pedestrians in the front and back right area, as well as the cars in the front left, front right and back area.

5 Conclusion

In this work, we have presented a novel radar-camera fusion architecture on the BEV plane. We propose two radar encoders and show that they can be integrated into several camera-based architectures that use BEV features. In our experiments, the proposed radar-camera fusion outperforms the camera-only baselines by a large margin, demonstrating its effectiveness. Without tuning the model for the test submission to avoid unrealistic computational cost, we outperform previously published radar-camera fusion networks. Our qualitative evaluation shows improved localization accuracy at daytime and higher recall at nighttime. In future work, we want to study the potential of BEV-based radar-camera fusion with the high-resolution, 3+1D radar sensors appearing in recently introduced datasets.

Acknowledgment

This work is partly funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) and partly financed by the European Union in the frame of NextGenerationEU within the project "Solutions and Technologies for Automated Driving in Town" (FKZ 19A22006P) and partly funded by the Federal Ministry of Education and Research Germany in the funding program Photonics Research Germany under the project FUMOS (FKZ 13N16302). The authors would like to thank the consortia for the successful cooperation.

References

1. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. arXiv preprint arXiv:1903.11027 (2019)
2. Chadwick, S., Maddern, W., Newman, P.: Distant vehicle detection using radar and vision. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 8311–8317. IEEE (2019)
3. Danzer, A., Griebel, T., Bach, M., Dietmayer, K.: 2d car detection in radar data with pointnets. In: IEEE Intelligent Transportation Systems Conference (ITSC). pp. 61–66. IEEE (2019)
4. Di Feng, Haase-Schutz, C., Rosenbaum, L., Hertlein, H., Glaeser, C., Timm, F., Wiesbeck, W., Dietmayer, K.: Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. IEEE Transactions on Intelligent Transportation Systems (2020)
5. Dreher, M., Erelik, E., Banziger, T., Knoll, A.: Radar-based 2d car detection using deep neural networks. In: IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC). pp. 1–8. IEEE (2020)
6. Drews, F., Di Feng, Faion, F., Rosenbaum, L., Ulrich, M., Glaser, C.: Deepfusion: A robust and modular 3d object detector for lidars, cameras and radars. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 560–567. IEEE (2022)
7. Fei, J., Chen, W., Heidenreich, P., Wirges, S., Stiller, C.: Semanticvoxels: Sequential fusion for 3d pedestrian detection using lidar point cloud and semantic segmentation. In: IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI). pp. 185–190. IEEE (2020)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
9. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)
10. Huang, J., Huang, G., Zhu, Z., Du, D.: Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. arXiv preprint arXiv:2112.11790 (2021)
11. Hung, W.C., Kretzschmar, H., Casser, V., Hwang, J.J., Anguelov, D.: Let-3d-ap: Longitudinal error tolerant 3d average precision for camera-only 3d detection. arXiv preprint arXiv:2206.07705 (2022)
12. Hwang, J.J., Kretzschmar, H., Manela, J., Rafferty, S., Armstrong-Crews, N., Chen, T., Anguelov, D.: Cramnet: Camera-radar fusion with ray-constrained cross-attention for robust 3d object detection. In: European Conference on Computer Vision. pp. 388–405. Springer (2022)
13. Kim, Y., Choi, J.W., Kum, D.: Grif net: Gated region of interest fusion network for robust 3d object detection from radar point cloud and monocular image. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 10857–10864. IEEE (2020)
14. Kim, Y., Kim, S., Choi, J.W., Kum, D.: Craft: Camera-radar 3d object detection with spatio-contextual fusion transformer. arXiv preprint arXiv:2209.06535 (2022)
15. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12697–12705 (2019)

16. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: Proceedings of the European conference on computer vision (ECCV). pp. 734–750 (2018)
17. Li, Y., Bao, H., Ge, Z., Yang, J., Sun, J., Li, Z.: Bevstereo: Enhancing depth estimation in multi-view 3d object detection with dynamic temporal stereo. arXiv preprint arXiv:2209.10248 (2022)
18. Li, Y., Ge, Z., Yu, G., Yang, J., Wang, Z., Shi, Y., Sun, J., Li, Z.: Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. arXiv preprint arXiv:2206.10092 (2022)
19. Liang, T., Xie, H., Yu, K., Xia, Z., Lin, Z., Wang, Y., Tang, T., Wang, B., Tang, Z.: Bevfusion: A simple and robust lidar-camera fusion framework. arXiv preprint arXiv:2205.13790 (2022)
20. Liu, Y., Wang, T., Zhang, X., Sun, J.: Petr: Position embedding transformation for multi-view 3d object detection. arXiv preprint arXiv:2203.05625 (2022)
21. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
22. Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. arXiv preprint arXiv:2205.13542 (2022)
23. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
24. Ma, X., Zhang, Y., Xu, D., Zhou, D., Yi, S., Li, H., Ouyang, W.: Delving into localization errors for monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4721–4730 (2021)
25. Major, B., Fontijne, D., Ansari, A., Teja Sukhavasi, R., Gowaikar, R., Hamilton, M., Lee, S., Grzechnik, S., Subramanian, S.: Vehicle detection with automotive radar using deep learning on range-azimuth-doppler tensors. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. p. 0 (2019)
26. Manhardt, F., Kehl, W., Gaidon, A.: Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2069–2078 (2019)
27. Meyer, M., Kusch, G.: Automotive radar dataset for deep learning based 3d object detection. In: 2019 16th European Radar Conference (EuRAD). pp. 129–132. IEEE (2019)
28. Mousavian, A., Anguelov, D., Flynn, J., Kosecka, J.: 3d bounding box estimation using deep learning and geometry. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7074–7082 (2017)
29. Nabati, R., Qi, H.: Rrpn: Radar region proposal network for object detection in autonomous vehicles. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 3093–3097. IEEE (2019)
30. Nabati, R., Qi, H.: Radar-camera sensor fusion for joint object detection and distance estimation in autonomous vehicles. arXiv preprint arXiv:2009.08428 (2020)
31. Nabati, R., Qi, H.: Centerfusion: Center-based radar and camera fusion for 3d object detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1527–1536 (2021)
32. Nobis, F., Geisslinger, M., Weber, M., Betz, J., Lienkamp, M.: A deep learning-based radar and camera sensor fusion architecture for object detection. In: 2019 Sensor Data Fusion: Trends, Solutions, Applications (SDF). pp. 1–7. IEEE (2019)

33. Palffy, A., Dong, J., Kooij, J.F.P., Gavrilu, D.M.: Cnn based road user detection using the 3d radar cube. *IEEE Robotics and Automation Letters* **5**(2), 1263–1270 (2020)
34. Palffy, A., Pool, E., Baratam, S., Kooij, J.F.P., Gavrilu, D.M.: Multi-class road user detection with 3+ 1d radar in the view-of-delft dataset. *IEEE Robotics and Automation Letters* **7**(2), 4961–4968 (2022)
35. Pang, S., Morris, D., Radha, H.: Transcar: Transformer-based camera-and-radar fusion for 3d object detection. *arXiv preprint arXiv:2305.00397* (2023)
36. Pillion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: *European Conference on Computer Vision*. pp. 194–210. Springer (2020)
37. Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum pointnets for 3d object detection from rgb-d data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 918–927 (2018)
38. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 652–660 (2017)
39. Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. *arXiv preprint arXiv:1811.08188* (2018)
40. Scheiner, N., Kraus, F., Appenrodt, N., Dickmann, J., Sick, B.: Object detection for automotive radar point clouds—a comparison. *AI Perspectives* **3**(1), 1–23 (2021)
41. Scheiner, N., Schumann, O., Kraus, F., Appenrodt, N., Dickmann, J., Sick, B.: Off-the-shelf sensor vs. experimental radar—how much resolution is necessary in automotive radar classification? *arXiv preprint arXiv:2006.05485* (2020)
42. Schumann, O., Hahn, M., Dickmann, J., Wöhler, C.: Semantic segmentation on radar point clouds. In: *2018 21st International Conference on Information Fusion (FUSION)*. pp. 2179–2186. IEEE (2018)
43. Simonelli, A., Bulo, S.R., Porzi, L., Lopez-Antequera, M., Kotschieder, P.: Disentangling monocular 3d object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1991–1999 (2019)
44. Stacker, L., Heidenreich, P., Rambach, J., Stricker, D.: Fusion point pruning for optimized 2d object detection with radar-camera fusion. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 3087–3094 (2022)
45. Ulrich, M., Braun, S., Kohler, D., Niederlohner, D., Faion, F., Glaser, C., Blume, H.: Improved orientation estimation and detection with hybrid object detection networks for automotive radar. *arXiv preprint arXiv:2205.02111* (2022)
46. Vora, S., Lang, A.H., Helou, B., Beijbom, O.: Pointpainting: Sequential fusion for 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4604–4612 (2020)
47. Wang, C., Ma, C., Zhu, M., Yang, X.: Pointaugmenting: Cross-modal augmentation for 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11794–11803 (2021)
48. Wang, T., Zhu, X., Pang, J., Lin, D.: Fcos3d: Fully convolutional one-stage monocular 3d object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 913–922 (2021)
49. Wang, Y., Chao, W.L., Garg, D., Hariharan, B., Campbell, M., Weinberger, K.Q.: Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8445–8453 (2019)

50. Wang, Y., Guizilini, V.C., Zhang, T., Wang, Y., Zhao, H., Solomon, J.: Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In: Conference on Robot Learning. pp. 180–191. PMLR (2022)
51. Wang, Z., Jia, K.: Frustum convnet: Sliding frustums to aggregate local point-wise features for amodal 3d object detection. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1742–1749. IEEE (2019)
52. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. *Sensors* **18**(10), 3337 (2018)
53. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV). pp. 767–783 (2018)
54. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11784–11793 (2021)
55. Zheng, L., Ma, Z., Zhu, X., Tan, B., Li, S., Long, K., Sun, W., Chen, S., Zhang, L., Wan, M.: Tj4dradset: A 4d radar dataset for autonomous driving. arXiv preprint arXiv:2204.13483 (2022)
56. Zhou, H., Ge, Z., Li, Z., Zhang, X.: Matrixvt: Efficient multi-camera to bev transformation for 3d perception. arXiv preprint arXiv:2211.10593 (2022)
57. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019)
58. Zhou, Y., Liu, L., Zhao, H., López-Benítez, M., Yu, L., Yue, Y.: Towards deep radar perception for autonomous driving: Datasets, methods, and challenges. *Sensors* **22**(11), 4208 (2022)