

# CaFNet: A Confidence-Driven Framework for Radar Camera Depth Estimation

Huawei Sun<sup>1,2</sup>, Hao Feng<sup>2</sup>, Julius Ott<sup>1,2</sup>, Lorenzo Servadei<sup>2</sup>, Robert Wille<sup>2</sup>

<sup>1</sup>Infineon Technologies AG, Neubiberg, Germany

<sup>2</sup>Technical University of Munich, Munich, Germany

E-mail: {huawei.sun, julius.ott}@infineon.com  
{hao.feng, lorenzo.servadei, robert.wille}@tum.de

**Abstract**—Depth estimation is critical in autonomous driving for interpreting 3D scenes accurately. Recently, radar-camera depth estimation has become of sufficient interest due to the robustness and low-cost properties of radar. Thus, this paper introduces a two-stage, end-to-end trainable Confidence-aware Fusion Net (CaFNet) for dense depth estimation, combining RGB imagery with sparse and noisy radar point cloud data. The first stage addresses radar-specific challenges, such as ambiguous elevation and noisy measurements, by predicting a radar confidence map and a preliminary coarse depth map. A novel approach is presented for generating the ground truth for the confidence map, which involves associating each radar point with its corresponding object to identify potential projection surfaces. These maps, together with the initial radar input, are processed by a second encoder. For the final depth estimation, we innovate a confidence-aware gated fusion mechanism to integrate radar and image features effectively, thereby enhancing the reliability of the depth map by filtering out radar noise. Our methodology, evaluated on the nuScenes dataset, demonstrates superior performance, improving upon the current leading model by 3.2% in Mean Absolute Error (MAE) and 2.7% in Root Mean Square Error (RMSE). Code: <https://github.com/harborsarah/CaFNet>

## I. INTRODUCTION

Dense depth estimation is an important task in autonomous driving, which helps to understand the 3D geometry of outdoor scenes. Following the first learning-based monocular depth estimation model [1], there have been studies enhancing task performance based on deep learning techniques [2]–[7]. Despite these advancements, the efficacy of monocular image-based methods is inherently constrained by the lack of depth cues in RGB images. To address this, numerous studies [8]–[10] have incorporated lidar data alongside images for depth completion tasks, yielding higher-quality depth maps.

While integrating lidar data offers improved scene understanding, lidar sensors are known to be sensitive to lighting and weather conditions. In contrast, radar sensors, present a cost-effective and weather-adaptable alternative. Moreover, radar sensors capture the velocity of moving objects, yielding more accurate object detection [11]–[13]. With the release of autonomous driving datasets [14, 15] that include radar sensors, the domain of camera-radar fusion is gaining traction. However, radar point clouds are notably sparser compared to lidar point clouds. Additionally, commonly used radar sensors often lack elevation resolution due to limited antennas along the elevation axis, leading to the absence of height

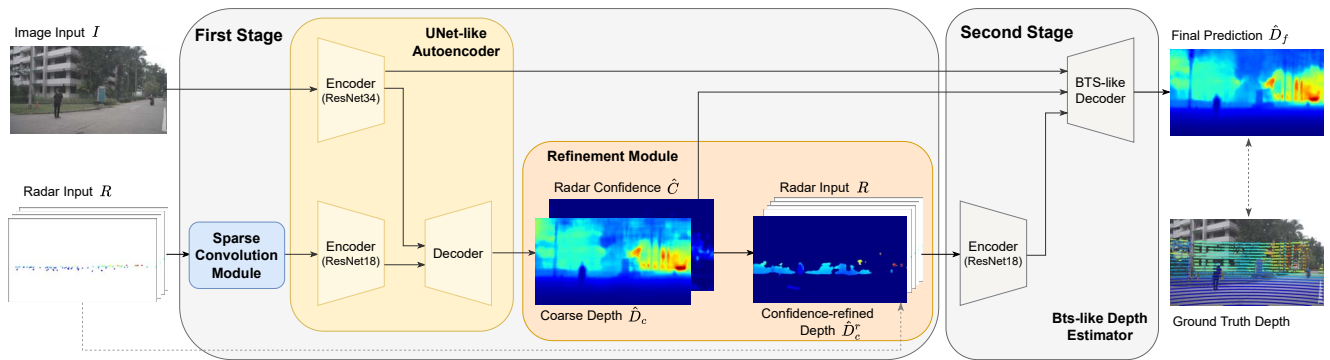
information in radar points. Moreover, the phenomenon of the radar multipath effect further complicates this scenario by introducing numerous ghost targets.



**Fig. 1:** Radar confidence map comparison. The ground truth is generated by our method, comparing the depth value of the radar point with the ground truth depth map within a selective region.

Given these challenges, effectively leveraging radar points while mitigating noise is a crucial concern. Several studies [11, 12, 16, 17] attempt to solve the sparsity and ambiguous height problems by extending radar points along the elevation axis, bringing additional errors into the radar data. In an alternative approach, RC-PDA [18] and RadarNet [19] employ two-stage networks that initially estimate radar confidence scores to generate semi-dense depth maps for further training. However, a predefined region with  $h \times w$  is selected for each radar point. This treats each radar point uniformly and does not consider the association between the radar point and its corresponding object in the 3D space.

To overcome the obstacles, we propose CaFNet, a novel two-stage, end-to-end trainable network for depth estimation. Our method incorporates a Sparse Convolution Module (SCM), employing sparse convolutional layers [20] to handle radar data sparsity before being processed by the encoder. In the first stage, a radar confidence map along with a coarse depth map is estimated, which guides the training of the second stage. To better guide the confidence map, we propose a ground truth generation method that considers the association between the radar points and the detected objects. An example of the confidence maps is illustrated in Fig. 1. The key innovation in our work is introducing Confidence-aware Gated Fusion (CaGF) in the second stage. This fusion technique integrates radar and image features by taking into account the confidence scores of individual radar pixels. Our model leverages a single radar scan and a single RGB image for depth estimation. The CaFNet outperforms existing solutions, establishing new benchmarks



**Fig. 2:** CaFNet model architecture, including two end-to-end trainable phases. Initially, the first stage focuses on estimating a radar confidence map and a coarse depth map. Subsequently, through the refinement module, a confidence-enhanced depth map is merged with the original radar input and then forwarded to a secondary radar encoder. The extracted radar features and image features, which, alongside the predicted confidence map, are input into the final decoder to generate the dense depth map.

on the nuScenes dataset [14].

To summarize, our contributions are as follows:

- We introduce the Confidence-aware Fusion Net (CaFNet), which surpasses previous state-of-the-art methods in radar-camera depth estimation on the nuScenes dataset.
- A novel approach for generating the ground truth of the radar confidence map is presented, enhancing the reliability of guiding the confidence generation.
- The Confidence-aware Gated Fusion (CaGF) technique is proposed, effectively reducing the propagation of erroneous data, thereby improving the overall performance of the depth estimation.

## II. RELATED WORK

### A. Monocular Depth Estimation

Monocular depth estimation is the process of determining the depth of each pixel in an image using a single RGB image. The majority of current methods employ encoder-decoder architectures to derive depth maps from input images. A significant advancement in this field was made by Eigen *et al.* [21], who introduced the use of Convolutional Neural Networks (CNNs) in a multi-scale network for depth estimation. This included the development of a scale-invariant loss function to address the issue of scale ambiguity. DORN [22] first reinterpreted depth regression as a classification task, focusing on predicting depth ranges instead of exact values with predefined bin centers. Subsequent studies, such as AdaBins [4] and BinsFormer [23], enhanced this approach by learning to predict more accurate bin centers, addressing the variation in depth distribution across different frames. Incorporating geometric priors has also been shown to be effective. For instance, GeoNet [5] uses a geometric network to simultaneously infer surface normals and depth maps that are geometrically consistent. As proposed in BTS [3], the local planar assumption guides the upsampling modules in the decoding phase. P3Depth [24] introduces an approach that selectively leverages information from coplanar pixels, using offset vector fields and a mean plane loss to refine depth predictions. Another line of research focuses on refining coarse depth maps through

spatial propagation networks, initially proposed in SPN [25]. CSPN [2] improved this by using convolutional layers to learn local affinities, which better connect each pixel with its neighboring pixels.

### B. Camera and Radar-based Depth Estimation

Camera radar fusion-based techniques combine sparse radar point clouds with camera images for depth map prediction. These methods present more challenges than lidar-camera fusion due to the inherent sparsity and noise in radar point clouds. Lin *et al.* [26] developed a two-stage encoder-decoder network specifically designed to filter out noisy radar points, using lidar points as a reference. An extension of this approach by [16] involves augmenting each radar point's height value to produce a denser radar projection, which is then integrated with camera images using a regression network based on DORN's [22] methodology. RC-PDA [18] tackles the uncertainty of projecting radar points onto the image plane by first learning a radar-to-pixel association. Once the association is established, a second network estimates depth. These studies typically use multiple radar sweeps to enhance the density of the point cloud. Singh *et al.* [19] propose a two-stage network, where the first stage learns a one-to-many mapping, generating quasi-dense radar depth maps, which are then processed in the second stage for final depth estimation. However, a significant drawback of the two-stage networks in RC-PDA [18] and RadarNet [19] is their separate training phases. This necessitates local storage of intermediate data post the initial training stage, resulting in a time-intensive process that challenges real-time prediction. Moreover, a predefined number of radar points are selected in RadarNet [19], with the empirical evidence suggesting a mere four points per frame. This severely constrains the utility of radar data.

## III. APPROACH

This section introduces the principal innovations of our work. Initially, we detail the comprehensive architecture of our CaFNet in Sec. III-A. Subsequently, Sec. III-B describes our novel approach for generating the ground truth of radar confidence maps. Following this, the refinement module is

described in Sec. III-C. Finally, we present the decoder, incorporating our newly proposed CaGF technique.

### A. Model Architecture

As illustrated in Fig. 2, our CaFNet processes an RGB image  $I \in \mathbb{R}^{H \times W \times 3}$  and a radar-projected image  $R \in \mathbb{R}^{H \times W \times C_R}$  as inputs. Typically, the ground truth depth  $D_{gt} \in \mathbb{R}^{H \times W}$  generated by lidar is sparse. Thus, for the supervision, following [19], we accumulate point clouds by projecting from neighboring frames to form a denser ground truth map  $D_{acc} \in \mathbb{R}^{H \times W}$ . The radar-projected image is generated by directly projecting every radar point onto the image plane, carrying specific information, for example, velocities and Radar Cross Section (RCS). In the first stage, feature extraction from the RGB image is conducted using a ResNet-34 [27] backbone. In contrast, the radar input undergoes initial refinement through the SCM, which stacks four sparse convolution layers [20], before being encoded by a ResNet-18 backbone. These multi-scale image and radar features are then concatenated and fed into a UNet decoder [28], which leverages skip connections to efficiently transmit information from the encoder to the decoder. The decoder outputs two maps. The first output is a coarse depth map  $\hat{D}_c \in \mathbb{R}^{H \times W}$ , which is trained against the ground truth depth map  $D_{acc} \in \mathbb{R}^{H \times W}$ . The second is a radar confidence map  $\hat{C} \in [0, 1]^{H \times W}$ , which signifies the probable correspondence between projected radar points and image pixels and is trained against a binary ground truth confidence map  $C \in \{0, 1\}^{H \times W}$ , detailed in Sec. III-B.

The refinement module, utilizing  $\hat{C}$  and  $\hat{D}_c$  as inputs, produces a confidence-refined depth map  $\hat{D}_c^r \in \mathbb{R}^{H \times W}$ . This map is concatenated with  $R$ , yielding a new radar input  $R' \in \mathbb{R}^{H \times W \times (C_R + 1)}$  and forwarded to the second stage. Sec. III-C provides a more detailed explanation.

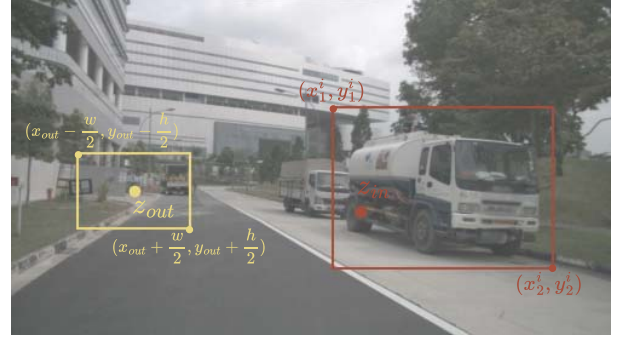
During the second stage,  $R'$  is encoded via a separate ResNet-18 backbone. Subsequently, a BTS-like depth estimator [3] is employed, which takes the newly extracted radar features, the image features, and the predicted radar confidence as inputs. Throughout this stage, features from different layers are fused using a CaGF mechanism. Leveraging multi-scale features, a final depth map  $\hat{D}_f \in \mathbb{R}^{H \times W}$  is estimated. Further details can be found in Sec. III-D.

### B. Radar Confidence Map

In the first stage in [19], a set of  $K$  radar points is chosen as input. For each point,  $K$  ground truth confidence maps of size  $288 \times 900$  are generated. However, with its substantially large predefined area, this approach results in significant computational demands. Moreover, since their model cannot handle various numbers of points as input, they select only four radar points during training, which leads to inefficiency.

In contrast, our method diverges by incorporating all radar points in each frame. We also pay particular attention to the relationship between radar points and the associated objects to more accurately determine probable radar projection areas. Specifically, for a frame, we work with a radar point cloud comprising  $M$  points,  $\mathbf{z} = \{z_m\}_{m=1}^M$ , and an accumulated ground truth depth map,  $D_{acc}$ . Additionally, we identify  $\mathcal{B} =$

$\{B^1, B^2, \dots, B^n\}$  as a set of object bounding boxes within the frame, where  $n$  is the number of detected objects.



**Fig. 3:** Radar confidence selective region. **red:** radar point is located in an object, **yellow:** radar point is not associated with any object.

We categorize the radar points into two groups. The first,  $\mathbf{z}_{in} = \{z_{in,p}\}_{p=1}^P$ , includes points inside any 3D bounding box, where  $z_{in} \in B^i$  signifies that the point is in the  $i^{th}$  box. The second group,  $\mathbf{z}_{out} = \{z_{out,q}\}_{q=1}^Q$ , contains points not associated with any objects, typically regarded as noise. Fig. 3 depicts these point types and their specific regions.

For a point  $z_{in}$  at pixel coordinates  $(x_{in}, y_{in})$  linked to bounding box  $B^i$ , the selected region  $\mathcal{R}_{in}$  is defined as:

$$\mathcal{R}_{in} = \{(x, y) \mid x_1^i \leq x \leq x_2^i, y_1^i \leq y \leq y_2^i\}, \quad (1)$$

where  $(x_1^i, y_1^i)$  and  $(x_2^i, y_2^i)$  represent the top-left and bottom-right corners of  $B^i$ , respectively.

Conversely, for a point  $z_{out}$  at coordinates  $(x_{out}, y_{out})$ , the region  $\mathcal{R}_{out}$  is selected using a pre-defined patch size  $(w, h)$ :

$$\mathcal{R}_{out} = \{(x, y) \mid x_{out} - \frac{w}{2} \leq x \leq x_{out} + \frac{w}{2}, y_{out} - \frac{h}{2} \leq y \leq y_{out} + \frac{h}{2}\}. \quad (2)$$

The ground truth confidence is generated by comparing the absolute depth difference of the point's depth with  $D_{acc}$  within the selective region. For the  $m^{th}$  radar point  $z_m$  with depth  $d_m$ , the confidence value of each pixel within its selective region  $\mathcal{R}_m$ :

$$C(i, j) = \begin{cases} 1 & \text{if } |D_{acc}(i, j) - d_m| \leq \tau \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Here,  $D_{acc}(i, j)$  represents the ground truth depth of pixel  $(i, j)$  and  $\tau$  is the tolerance threshold, chosen as 0.4 meters.

### C. Refinement Module

To facilitate the learning of information from  $\hat{D}_c$  and  $\hat{C}$  by the second radar encoder, our initial concept involved concatenating these two maps with  $R$ . However, due to the sparsity property of  $R$ , the detailed information from  $\hat{D}_c$  might overwhelm the encoder with excessive information. To address this, we introduce a refinement module designed to streamline the information content from  $\hat{D}_c$ , utilizing the predicted radar confidence map as guidance.

In this refinement module, we generate a radar confidence mask  $\hat{M} \in \mathbb{R}^{H \times W}$  by setting a predefined threshold  $T$ :

$$\hat{M}(i, j) = \begin{cases} 1 & \text{if } \hat{C}(i, j) \geq T \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Followed by this, the confidence-refined depth  $\hat{D}_c^r$  is calculated by element-wise multiplication between  $\hat{M}$  and  $\hat{D}_c$

$$\hat{D}_c^r = \hat{M} \odot \hat{D}_c. \quad (5)$$

After the refinement process,  $\hat{D}_c^r$  is concatenated with  $R$ , creating a new input  $R'$  and then fed into the second stage.

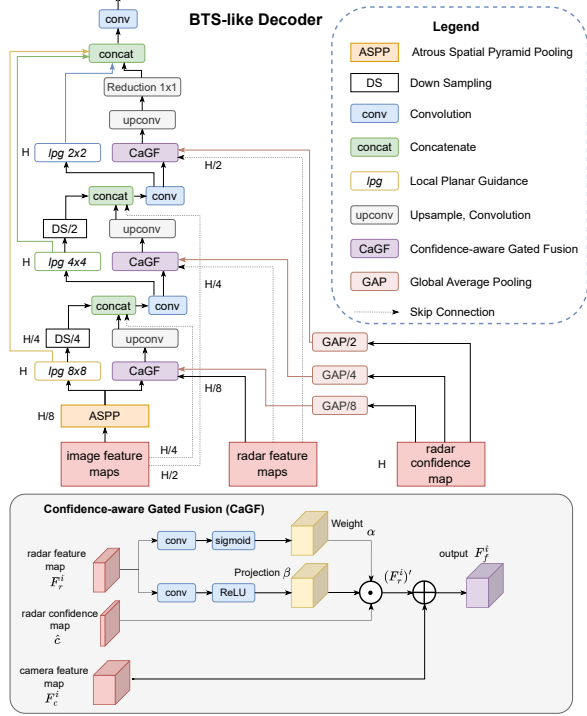


Fig. 4: BTS-like Decoder.

#### D. Bts-like Depth Estimator

Our final depth estimation model adopts the framework of BTS [3]. To enhance its performance, we introduce the CaGF—a novel mechanism designed to integrate image and radar features. This process incorporates  $\hat{C}$ , enabling the selective filtration of noisy radar data. The decoder’s architecture is depicted in Fig. 4.

**Confidence-aware Gated Fusion:** At the  $i^{\text{th}}$  stage, the radar feature  $F_r^i$  and the image feature  $F_c^i$  are in shape of  $\frac{H}{2^i} \times \frac{W}{2^i} \times C_r$  and  $\frac{H}{2^i} \times \frac{W}{2^i} \times C_c$ , respectively. Building upon the approach in [19], we derive a set of weights  $\alpha = \sigma(p^T F_r^i) \in [0, 1]$  and projections  $\beta = q^T F_r^i$  from  $F_r^i$ . Here,  $p$  and  $q$  denote trainable parameters, and  $\sigma$  symbolizes the sigmoid activation function, ensuring  $\alpha$  scales between 0 and 1. Additionally, we process the predicted radar confidence  $\hat{C} \in [0, 1]^{H \times W}$  to obtain a resized confidence  $\hat{c}$  by employing Global Average Pooling (GAP) with a stride of  $2^i$  on  $\hat{C}$ .

The enhancement of the radar feature is formulated as  $(F_r^i)' = \alpha \cdot \beta \cdot \hat{c}$ . This refinement step leverages the radar confidence to more effectively distinguish the significance of various positions, a critical consideration given that the

majority of radar input pixels are typically zero-valued. Consequently, the composite feature for this stage is obtained as  $F_f^i = (F_r^i)' + F_c^i$ , integrating both radar and image features for a more nuanced representation.

#### E. Loss Functions

Our model employs an end-to-end training strategy with several loss functions. For depth estimation, we use an L1 loss on the predictions  $\hat{D}_c$  and  $\hat{D}_f$  from both the coarse and final stages, respectively. The depth loss is given by:

$$L_{Depth} = \frac{m}{|\Omega|} \sum_{x \in \Omega} |D_{acc}(x) - \hat{D}_c(x)| + \frac{1}{|\Omega|} \sum_{x \in \Omega} |D_{acc}(x) - \hat{D}_f(x)|, \quad (6)$$

where  $m$  is a weighting factor, chosen as 0.5, and  $\Omega$  represents the set of pixels where  $D_{acc}$  is valid.

Additionally, an edge-aware smoothness loss is applied to the final depth map,  $\hat{D}_f$ , follows [26], formulated as:

$$L_{Smooth} = |\nabla_u \hat{D}_f| e^{-|\nabla_u(I)|} + |\nabla_v \hat{D}_f| e^{-|\nabla_v(I)|}, \quad (7)$$

where  $\nabla_u$  and  $\nabla_v$  are the horizontal and vertical image gradients, respectively.

For radar confidence, we minimize a binary cross-entropy loss between the ground truth  $C$  and the predicted map  $\hat{C}$ :

$$L_{Conf} = -\frac{1}{|\Omega_C|} \sum_{x \in \Omega_C} (C(x) \log(\hat{C}(x)) + (1 - C(x)) \log(1 - \hat{C}(x))), \quad (8)$$

where  $\Omega_C$  denotes the image domain.

The final loss  $L$  is a weighted sum of the individual losses  $L = L_{Depth} + L_{Conf} + \lambda L_{Smooth}$ . Here,  $\lambda$  is a hyperparameter, chosen as  $1e^{-3}$  based on empirical results.

## IV. EXPERIMENTS

This section first describes the dataset and implementation setups. Then, we introduce the evaluation metrics and compare our method with existing approaches in qualitative and quantitative ways. Afterwards, we show the efficacy of the radar confidence map. Finally, we conducted ablation studies to underscore the robustness of the proposed techniques.

#### A. Dataset and Implementation Details

We utilize nuScenes [14], a leading-edge dataset for autonomous driving, to train and test our model’s performance. This dataset is captured via multiple sensors, including cameras, radars, and lidar, mounted on a vehicle navigating through Boston and Singapore. Here, we utilize 700 scenes for training, 150 for validation, and 150 for testing.

During training, we follow [19] to accumulate 80 future and 80 past lidar frames by projecting the lidar point cloud at each frame to the current frame, thereby creating a denser depth map. Subsequently, we apply scaffolding [31] to derive the interpolated depth map  $D_{acc}$ . During evaluation, we use the single-frame lidar depth map  $D_{gt}$  provided by [14].

Our model is implemented using PyTorch [32] and trained on a Nvidia<sup>®</sup> Tesla A30 GPU. We adopt the Adam optimizer,

**TABLE I:** Performance Comparison on nuScenes Official Test Set.

Max Eval Distance	Method	Sensors		Metrics							
		Image	Radar	MAE ↓	RMSE ↓	AbsRel ↓	log10 ↓	RMSElog ↓	$\delta_1$ ↑	$\delta_2$ ↑	$\delta_3$ ↑
50m	BTS [3]	✓		1.937	3.885	0.116	0.045	0.179	0.883	0.957	1.937
	RC-PDA [18]	✓	✓	2.225	4.159	0.106	0.051	0.186	0.864	0.944	0.974
	RC-PDA-HG [18]	✓	✓	2.210	4.234	0.121	0.052	0.194	0.850	0.942	0.975
	DORN [16]	✓	✓	1.898	3.928	0.100	0.050	<b>0.164</b>	0.905	0.962	0.982
	RadarNet [19]	✓	✓	1.706	3.742	0.103	0.041	0.170	0.903	<b>0.965</b>	<b>0.983</b>
	CaFNet (Ours)	✓	✓	<b>1.674</b>	<b>3.674</b>	<b>0.098</b>	<b>0.038</b>	<b>0.164</b>	<b>0.906</b>	0.963	<b>0.983</b>
70m	BTS [3]	✓		2.346	4.811	0.119	0.047	0.188	0.872	0.952	0.979
	RC-PDA [18]	✓	✓	3.338	6.653	0.122	0.060	0.225	0.822	0.923	0.965
	RC-PDA-HG [18]	✓	✓	3.514	7.070	0.127	0.062	0.235	0.812	0.914	0.960
	DORN [16]	✓	✓	2.170	4.532	0.105	0.055	<b>0.170</b>	0.896	0.960	0.982
	RadarNet [19]	✓	✓	2.073	4.591	0.105	0.043	0.181	0.896	<b>0.962</b>	0.981
	CaFNet (Ours)	✓	✓	<b>2.010</b>	<b>4.493</b>	<b>0.101</b>	<b>0.040</b>	0.174	<b>0.897</b>	0.961	<b>0.983</b>
80m	BTS [3]	✓		2.467	5.125	0.120	0.048	0.191	0.869	0.951	0.979
	AdaBins [4]	✓		3.541	5.885	0.197	0.089	0.261	0.642	0.929	0.977
	P3Depth [24]	✓		3.130	5.838	0.165	0.065	0.222	0.804	0.934	0.974
	LapDepth [29]	✓		2.544	5.151	0.117	0.049	0.187	0.865	0.953	0.981
	RC-PDA [18]	✓	✓	3.721	7.632	0.126	0.063	0.238	0.813	0.914	0.960
	RC-PDA-HG [18]	✓	✓	3.664	7.775	0.138	0.064	0.245	0.806	0.909	0.957
	DORN [16]	✓	✓	2.432	5.304	0.107	0.056	0.177	0.890	0.960	<b>0.981</b>
	RCDPT <sup>†</sup> [30]	✓	✓	-	5.165	0.095	-	-	0.901	<b>0.961</b>	<b>0.981</b>
	RadarNet [19]	✓	✓	2.179	4.899	0.106	0.044	0.184	0.894	0.959	0.980
	CaFNet (Ours)	✓	✓	<b>2.109</b>	<b>4.765</b>	<b>0.101</b>	<b>0.040</b>	<b>0.176</b>	<b>0.895</b>	0.959	<b>0.981</b>

<sup>†</sup> These results come from the paper that tests the model performance on a different test set, leading to the metrics being less comparable.

starting with an initial learning rate of  $1e^{-4}$  and applying a polynomial decay rate with a power of  $p = 0.9$ . The training process spans 150 epochs with batches of 10. To mitigate overfitting, image augmentation techniques such as random flipping and adjustments to contrast, brightness, and color are employed. Additionally, random cropping to  $352 \times 704$  pixels is performed during training to further enhance the model’s robustness.

### B. Evaluation Metrics

Table. II defines the evaluation metrics, where the predicted and ground truth depth maps are represented by  $\hat{d}$  and  $d_{gt}$ .  $\Omega$  denotes the set of pixels where  $d_{gt}$  is valid.

**TABLE II:** Metrics definition for depth estimation task.

	Definition
MAE	$\frac{1}{ \Omega } \sum_{x \in \Omega}  \hat{d}(x) - d_{gt}(x) $
RMSE	$(\frac{1}{ \Omega } \sum_{x \in \Omega}  \hat{d}(x) - d_{gt}(x) ^2)^{1/2}$
AbsRel	$\frac{1}{ \Omega } \sum_{x \in \Omega}  \hat{d}(x) - d_{gt}(x)  / d_{gt}(x)$
log10	$\frac{1}{ \Omega } \sum_{x \in \Omega}  \log_{10} \hat{d}(x) - \log_{10} d_{gt}(x) $
RMSElog	$\sqrt{\frac{1}{ \Omega } \sum_{x \in \Omega} \ \log_{10} \hat{d}(x) - \log_{10} d_{gt}(x)\ ^2}$
$\delta_n$ threshold	$\delta_n = \{ \hat{d}(x) : \max(\frac{\hat{d}(x)}{d_{gt}(x)}, \frac{d_{gt}(x)}{\hat{d}(x)}) < 1.25^n\} /  \Omega $

### C. Quantitative Results

We benchmark our method against existing image-only techniques [3, 4, 24, 29] and radar-camera fusion strategies [16, 18, 19, 30] on the official nuScenes test set, utilizing metrics in TABLE II. We conduct comparisons at 50, 70, and 80-meter intervals, with the results presented in TABLE I. Our model surpasses image-only depth estimators across all metrics. Specifically, comparing to our baseline [3], our approach exhibits an MAE reduction of 13.6%, 14.3%, and 14.5%, and RMSE decrease of 5.4%, 6.6%, and 7.0% at the respective distances. These improvements underscore the advantage of integrating radar data, which reliably captures long-range details that are often not visible in RGB imagery.

Following [19], our framework utilizes a single radar scan to ensure feasibility in real-world applications. In contrast,

techniques such as those presented in [16, 18, 26] leverage multiple radar sweeps to augment the dataset. Despite the limited radar data, our proposed technique demonstrates superior performance in nearly all evaluated metrics. Compared to the leading model, RadarNet [19], we achieve an improvement of 1.9%, 3.0%, and 3.2% in MAE for the stated ranges, further establishing the efficacy of our approach.

### D. Qualitative Results

Fig. 5 compares the performance of our method with the baseline [3] and the leading fusion model [19] across distances up to 80 meters under two scenarios. In the first scenario, our model precisely identifies the truck within the indicated regions (highlighted with red circles), unlike the other methods, which merge the truck’s top with the sky. Moreover, BTS fails to recognize the pedestrian, who is nearly visible on its depth map. The second row illustrates that under rainy conditions, the resulting blurred RGB image hinders the other models from delivering an accurate prediction. For instance, the upper part of the central vehicle can hardly be detected. In contrast, our method robustly differentiates and correctly measures the depth of each object. Additionally, depth maps produced by RadarNet often exhibit discontinuities, which can constrain the 3D reconstruction process. Our method, however, generates accurate depth maps that clearly distinguish individual objects.

### E. Efficacy of Radar Confidence

Fig. 6 showcases the radar confidence map with radar points marked in blue. The maps indicate the probable projection region for each radar point, aiming to mitigate the effects of noisy and inaccurate radar detections. Areas highlighted in red are identified as highly probable projection zones. Notably, as highlighted in yellow circles, a majority of the radar points are encompassed by regions with markedly low confidence scores. This observation indicates that these points are predominantly noise, which could lead to misleading during training. Involving this confidence map in

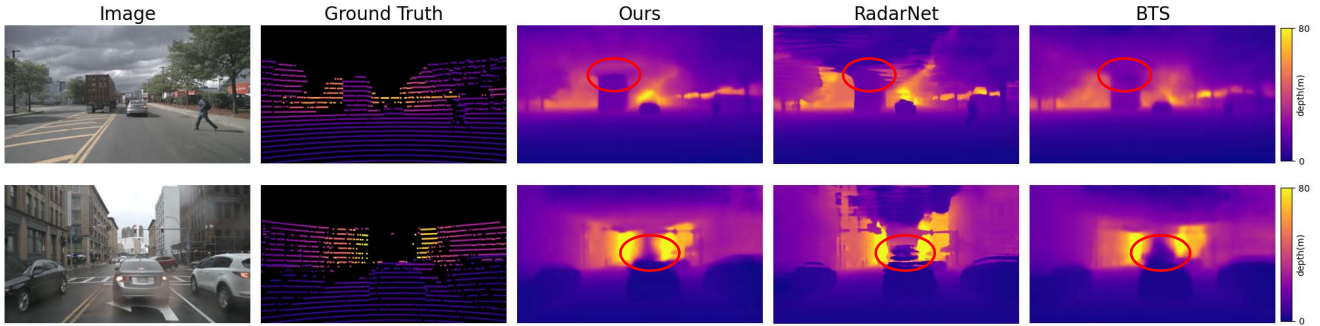


Fig. 5: Qualitative comparison on nuScenes test set. Column 1 shows the RGB image; column 2 plots the ground truth depth map. We compare our result with the RadarNet and our baseline BTS at 80 meters depth range.

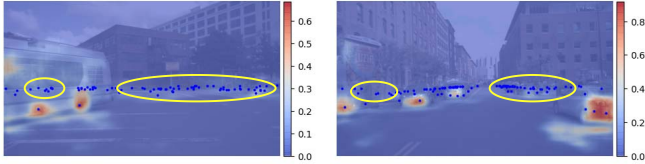
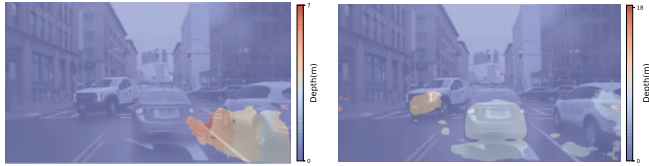


Fig. 6: Radar confidence map on top of the RGB image. The blue dots visualize the radar projected points. The confidence map indicates the possible projection surface of the radar points.



(a) Quasi-dense map estimated from the first stage of RadarNet [19]. (b) Generated refined depth after the refinement module.

Fig. 7: Qualitative comparison between the quasi-dense map [19] and our refined depth map.

the fusion process reduces the influence of such noisy data, thereby preserving and emphasizing valuable information.

Subsequently, the refinement module processes  $\hat{C}$  and  $\hat{D}_c$  to produce a refined depth map. We then evaluate our refined depth map against the quasi-dense map generated by RadarNet [19]. In RadarNet, the quasi-dense maps generated during the initial stage serve as the radar input for the subsequent stage. As depicted in Fig. 7, our refined depth map showcases significant depth details at locations corresponding to radar points with a confidence score exceeding 0.4. Conversely, the quasi-dense map from [19] includes data that lacks coherence and reliability, rendering it unsuitable as a dependable input for the final depth estimation process.

TABLE III: Ablation study about the radar confidence.

	MAE ↓	RMSE ↓	AbsRel ↓	$\delta_1$ ↑
GT from [19]	2.192	4.831	0.106	0.889
GT from [18]	2.196	4.838	0.106	0.890
w/o RM	2.189	4.844	0.105	0.890
CaGF → Add	2.201	4.855	0.108	0.891
CaGF → Concat	2.198	4.849	0.107	0.890
CaGF → GF	2.194	4.846	0.107	0.889
w/o confidence	2.208	4.878	0.109	0.890
Ours	<b>2.109</b>	<b>4.765</b>	<b>0.101</b>	<b>0.895</b>

#### F. Ablation Study

We conduct ablation studies to ascertain the impact of the proposed components. The results are detailed as follows:

**Effectiveness of radar confidence:** TABLE III showcases the experimental outcomes related to radar confidence. First, we benchmark our ground truth (GT) generation approach against the method proposed in [19] and [18]. Compared to [19], which defines a selective region of  $288 \times 900$ , our methodology yields a 3.8% improvement in MAE.

Secondly, in the Refinement Module (RM),  $\hat{C}$  is employed to generate the confidence-refined depth, concatenated with  $R$ , and progressed to the second stage. We evaluate this approach against a baseline that directly concatenates  $\hat{D}_c$ ,  $\hat{C}$ , and  $R$ , rather than refinement. The absence of RM leads to an overload of detailed information alongside sparse radar data, potentially misleading the model’s feature extraction.

Most importantly, we show the effectiveness of the CaGF by comparing it with the Gated Fusion (GF) [19] and the element-wise addition and concatenation. Our CaGF is able to distinguish the relevance of individual pixels within radar features during the fusion process, enhancing MAE by 3.9%.

Finally, to underscore the significance of the confidence map, we designed an additional experiment in which the first stage only predicts a coarse depth map  $\hat{D}_c$ . Consequently, within the RM,  $\hat{D}_c$  is concatenated with  $R$  and processed by the second encoder. In the decoding phase, the GF mechanism is employed to fuse the radar and image features. This setup aims to show the vital role of the confidence map in enhancing the accuracy and reliability of depth estimation.

TABLE IV: Ablation study about the SCM.

	MAE ↓	RMSE ↓	AbsRel ↓	$\delta_1$ ↑
w/o SCM	2.138	4.780	0.105	0.892
Ours (final)	<b>2.109</b>	<b>4.765</b>	<b>0.101</b>	<b>0.895</b>

**Effectiveness of SCM:** Table IV presents the outcomes demonstrating that the integration of the SCM leads to a 1.4% enhancement in MAE. This improvement underscores SCM’s efficacy in extracting useful features from sparse data.

## V. CONCLUSION

Depth estimation task using radar and camera data presents considerable challenges, primarily due to the inherent sparsity and ambiguity in radar data. In response to these challenges, we introduced CaFNet that operates in two stages. The first stage generates a radar confidence map and a preliminary depth map. In the second stage, we implement an innovative fusion technique that leverages the confidence scores to enhance the feature fusion process.

This method discriminates between the radar pixels based on their informational value, without incorporating additional misleading data, thereby significantly improving the accuracy of the depth maps. Our evaluations of the nuScenes dataset demonstrate the superiority of our approach over existing methodologies. Through our research, we provide a robust solution of radar-camera depth estimation, which improves the MAE by 3.2%. Moving forward, we aim to enhance the utilization of radar data by incorporating 3D geometric information more comprehensively.

## VI. ACKNOWLEDGEMENT

Research leading to these results has received funding from the EU ECSEL Joint Undertaking under grant agreement n° 101007326 (project AI4CSM) and from the partner national funding authorities the German Ministry of Education and Research (BMBF).

## REFERENCES

- [1] A. Saxena, S. Chung, and A. Ng, "Learning depth from single monocular images," *Advances in neural information processing systems*, vol. 18, 2005.
- [2] X. Cheng, P. Wang, and R. Yang, "Learning depth with convolutional spatial propagation network," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 10, pp. 2361–2379, 2019.
- [3] J. H. Lee, M.-K. Han, D. W. Ko, and I. H. Suh, "From big to small: Multi-scale local planar guidance for monocular depth estimation," *arXiv preprint arXiv:1907.10326*, 2019.
- [4] S. F. Bhat, I. Alhashim, and P. Wonka, "Adabins: Depth estimation using adaptive bins," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4009–4018.
- [5] X. Qi, R. Liao, Z. Liu, R. Urtasun, and J. Jia, "Geonet: Geometric neural network for joint depth and surface normal estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [6] W. Yuan, X. Gu, Z. Dai, S. Zhu, and P. Tan, "Neural window fully-connected crfs for monocular depth estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3916–3925.
- [7] J. Hu, M. Ozay, Y. Zhang, and T. Okatani, "Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries," in *2019 IEEE winter conference on applications of computer vision (WACV)*. IEEE, 2019, pp. 1043–1051.
- [8] F. Ma and S. Karaman, "Sparse-to-dense: Depth prediction from sparse depth samples and a single image," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4796–4803.
- [9] J. Park, K. Joo, Z. Hu, C.-K. Liu, and I. So Kweon, "Non-local spatial propagation network for depth completion," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*. Springer, 2020, pp. 120–136.
- [10] Y. Chen, B. Yang, M. Liang, and R. Urtasun, "Learning joint 2d-3d representations for depth completion," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 10 023–10 032.
- [11] 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.
- [12] H. Sun, H. Feng, G. Stettinger, L. Servadei, and R. Wille, "Multi-task cross-modality attention-fusion for 2d object detection," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 2023, pp. 3619–3626.
- [13] R. Nabati and H. 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*, 2021, pp. 1527–1536.
- [14] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [15] M. Bijelic, T. Gruber, F. Mannan, F. Kraus, W. Ritter, K. Dietmayer, and F. Heide, "Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather," in *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [16] C.-C. Lo and P. Vandewalle, "Depth estimation from monocular images and sparse radar using deep ordinal regression network," in *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2021, pp. 3343–3347.
- [17] H. Sun, H. Feng, G. Mauro, J. Ott, G. Stettinger, L. Servadei, and R. Wille, "Enhanced radar perception via multi-task learning: Towards refined data for sensor fusion applications," in *2024 IEEE Intelligent Vehicles Symposium (IV)*, 2024, pp. 3179–3184.
- [18] Y. Long, D. Morris, X. Liu, M. Castro, P. Chakravarty, and P. Narayanan, "Radar-camera pixel depth association for depth completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12 507–12 516.
- [19] A. D. Singh, Y. Ba, A. Sarker, H. Zhang, A. Kadambi, S. Soatto, M. Srivastava, and A. Wong, "Depth estimation from camera image and mmwave radar point cloud," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9275–9285.
- [20] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity invariant cnns," in *2017 international conference on 3D Vision (3DV)*. IEEE, 2017, pp. 11–20.
- [21] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in neural information processing systems*, vol. 27, 2014.
- [22] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao, "Deep ordinal regression network for monocular depth estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2002–2011.
- [23] Z. Li, X. Wang, X. Liu, and J. Jiang, "Binsformer: Revisiting adaptive bins for monocular depth estimation," *arXiv preprint arXiv:2204.00987*, 2022.
- [24] V. Patil, C. Sakaridis, A. Liniger, and L. Van Gool, "P3depth: Monocular depth estimation with a piecewise planarity prior," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1610–1621.
- [25] S. Liu, S. De Mello, J. Gu, G. Zhong, M.-H. Yang, and J. Kautz, "Learning affinity via spatial propagation networks," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [26] J.-T. Lin, D. Dai, and L. Van Gool, "Depth estimation from monocular images and sparse radar data," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 233–10 240.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [29] M. Song, S. Lim, and W. Kim, "Monocular depth estimation using laplacian pyramid-based depth residuals," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, no. 11, pp. 4381–4393, 2021.
- [30] C.-C. Lo and P. Vandewalle, "Rcdpt: Radar-camera fusion dense prediction transformer," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [31] A. Wong, X. Fei, S. Tsuei, and S. Soatto, "Unsupervised depth completion from visual inertial odometry," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1899–1906, 2020.
- [32] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," *ArXiv*, vol. abs/1912.01703, 2019.