

# RadarCam-Depth: Radar-Camera Fusion for Depth Estimation with Learned Metric Scale

Han Li<sup>1,†</sup>, Yukai Ma<sup>1,†</sup>, Yaqing Gu<sup>1</sup>, Kewei Hu<sup>1</sup>, Yong Liu<sup>1,\*</sup>, Xingxing Zuo<sup>2,\*</sup>

**Abstract**—We present a novel approach for metric dense depth estimation based on the fusion of a single-view image and a sparse, noisy Radar point cloud. The direct fusion of heterogeneous Radar and image data, or their encodings, tends to yield dense depth maps with significant artifacts, blurred boundaries, and suboptimal accuracy. To circumvent this issue, we learn to augment versatile and robust monocular depth prediction with the dense metric scale induced from sparse and noisy Radar data. We propose a Radar-Camera framework for highly accurate and fine-detailed dense depth estimation with four stages, including monocular depth prediction, global scale alignment of monocular depth with sparse Radar points, quasi-dense scale estimation through learning the association between Radar points and image patches, and local scale refinement of dense depth using a scale map learner. Our proposed method significantly outperforms the state-of-the-art Radar-Camera depth estimation methods by reducing the mean absolute error (MAE) of depth estimation by 25.6% and 40.2% on the challenging nuScenes dataset and our self-collected ZJU-4DRadarCam dataset, respectively. Our code and dataset will be released at <https://github.com/MOCKING/RadarCam-Depth>.

## I. INTRODUCTION

Perceiving the environment is critically important for autonomous driving, where accurate depth estimation is fundamental for dense reconstruction, 3D detection, and obstacle avoidance. Cameras and range sensors have been widely used for perceiving dense depth. Learned monocular depth (mono-depth) estimation methods based on CNN networks [1–5] have been prevalent in recent years due to their versatile applicability and plausible accuracy. They benefit from the solid contextual priors from extensive training on diverse datasets. While mono-depth networks excel in estimating up-to-scale depth, they fail to predict the accurate metric scale of depth. This limit arises from the inherent challenge of capturing scale with single-view cameras and the difficulty of learning the diverse scale in complex scenarios.

Range sensors like LiDAR and Radar can provide metric scale information of the scene [6, 7]. While LiDAR is renowned for its ability to generate dense and accurate point clouds, its widespread deployment faces challenges due to high costs, power consumption, and data bandwidth limits. In contrast, 3D Radar has witnessed remarkable advancements,

<sup>1</sup>The authors are with the Institute of Cyber-Systems and Control, Zhejiang University, Hangzhou, China.

<sup>2</sup>The author is with the Department of Computing and Mathematical Sciences, California Institute of Technology, USA.

\* Xingxing Zuo and Yong Liu are the corresponding authors (Email: [zuox@caltech.edu](mailto:zuox@caltech.edu); [yongliu@iipc.zju.edu.cn](mailto:yongliu@iipc.zju.edu.cn)).

<sup>†</sup> These authors contributed equally to this work.

This work is supported by NSFC 62088101 Autonomous Intelligent Unmanned Systems.



Fig. 1: Top: 3D visualization of the metric depth estimation from our proposed RadarCam-Depth; Middle: Our metric depth estimation overlaid on corresponding error map; Bottom: Error map of Mono-depth after scale-aligned to Radar points. Our depth estimation exhibits exceptional metric accuracy and fine details.

making it attractive in autonomous driving, owing to its affordability, low power consumption, and high resilience in challenging fog and smoke scenarios. The emerging 4D Radar additionally provides an elevation dimension with extended applicability. Fusing data from a single camera and a Radar for metric dense depth estimation becomes a promising research avenue [8–13]. It holds substantial significance in autonomous driving since its appealing characteristics, like cost-effectiveness, complementarity in sensing capabilities, and remarkable robustness and reliability.

However, sparsity, substantial noise in Radar data, and the imperfect cross-modal association between Radar points and image pixels pose challenges for dense depth estimation. Previous Radar-Camera methods treat the dense depth estimation as a depth completion problem [10, 12]. In these methods, the initial step involves associating the Radar depth to the camera pixels, generating a sparse or semi-dense depth map, which is then completed by an Unet-like network with a fusion of the Radar depth and image data. In this paper, we propose a novel paradigm, RadarCam-Depth, which capitalizes on robust and versatile scaleless monocular depth prediction and learns to assign metric dense scales to the mono-depth with Radar data. Our novel paradigm offers two main benefits: (i) We circumvent the direct fusion of raw data or encodings of heterogeneous Radar and camera data, thereby preventing aliasing artifacts and preserving high-fidelity fine details in dense depth estimation (see Fig.1). (ii) Unlike learning the depth completion with a wide con-

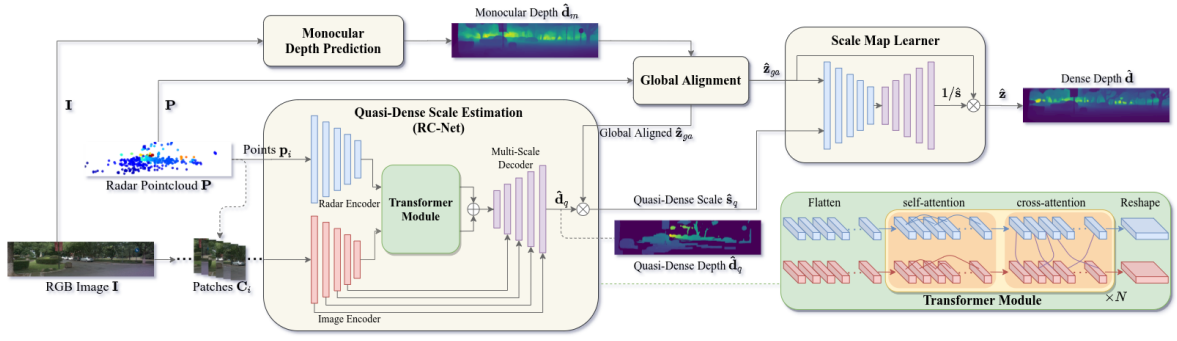


Fig. 2: The overall framework of our proposed RadarCam-Depth, comprised with four stages: monocular depth prediction, global alignment of mono-depth with sparse Radar depth, learned quasi-dense scale estimation, and scale map learner for refining local scale.  $\mathbf{d}$  and  $\mathbf{s}$  denotes the depth and scale, while  $\mathbf{z} = 1/\mathbf{d}$  is the inverse depth.

vergence basin, we essentially learn to complete the sparse scale obtained by aligning Radar depth with the mono-depth, which is more accessible and conducive to effective learning.

The primary contributions of this work are as follows: (i) We introduce the first approach that enhances the highly generalizable, scaleless mono-depth prediction with the dense metric scale intricately inferred from the noisy and sparse Radar data. (ii) We present a novel metric dense depth estimation framework that effectively fuses heterogeneous Radar and camera data. Our framework comprises four stages: mono-depth prediction, global scale alignment of the monocular depth, Radar-Camera quasi-dense scale estimation, and scale map learner for refining the quasi-dense scale locally. (iii) The proposed method is extensively tested on the nuScenes benchmark and our self-collected ZJU-4DRadarCam dataset. It outperforms the state-of-the-art (SOTA) techniques, substantially enhancing Radar-Camera dense depth estimation with high metric accuracy and strong generalizability. (iv) To fertilize future research in robust depth estimation, we will release our code and high-quality ZJU-4DRadarCam dataset, including raw 4D Radar data, RGB images, and meticulously generated ground truth from LiDAR measurements.

## II. RELATED WORK

### A. Monocular Depth Estimation

Monocular depth estimation is a challenging task due to the inherent scale ambiguity. Many researchers have tried to address this issue by integrating it with optical flow [14], uncertainty estimation [15], semantic segmentation [16], instance segmentation [17] and visual odometry [18]. Although some previous studies [3, 4, 19, 20] have achieved promising results in affine-invariant scaleless depth estimation across diverse datasets, recovering the metric scale remains a significant challenge. Some existing methods rely on inertial data to provide scale. To enhance the generalization, VI-SLAM [21] warps the input image to match the orientation prevailing in the training dataset. CodeVIO [22] proposes a tightly coupled VIO system with optimizable learned dense depth. It jointly estimates VIO poses and optimizes the predicted and encoded dense depth of specific keyframes efficiently. Xie et al. [23] utilize a flow-to-depth layer to

refine camera poses and generate depth proposals. They solve a multi-frame triangulation problem to enhance the estimation accuracy. Recently, Wofk et al. [5] introduced a framework for metric dense depth estimation from the VIO sparse depth and monocular depth prediction, which inspires our work. They first globally align the scaleless mono-depth with the metric VIO sparse depth and then learn to refine the dense scale of the globally aligned mono-depth.

### B. Depth Estimation from Radar-Camera Fusion

The fusion of Radar and camera data for metric depth estimation is an active research topic. Lin et al. [8] introduce a two-stage CNN-based pipeline that combines Radar and camera inputs to denoise Radar signals and estimate dense depth. Long et al. [9] propose a Radar-2-Pixel (R2P) network that utilizes radial Doppler velocity and induced optical flow from images to associate Radar points with corresponding pixel regions, enabling the synthesis of full-velocity information. They also achieve image-guided depth completion using Radar and video data [10]. Another approach, DORN [11] proposed by Lo et al., extends Radar points in the elevation dimension and applies deep ordinal regression network-based [24] feature fusion. Unlike other methods, R4dyn [13] creatively incorporates Radar as a weakly supervised signal into a self-supervised framework and employs Radar as an additional input to enhance the robustness. However, their method primarily focuses on vehicle targets and does not fully correlate all Radar points with a larger image area, resulting in lower depth accuracy. Recently, Singh et al. [12] present a method that relies solely on a single image frame and Radar point cloud. Their first-stage network infers the confidence scores of Radar-Pixel correspondence, generating a semi-dense depth map. They further employ a gated fusion network to control the fusion of multi-modal Radar-Camera data and predict the dense depth. However, all the above methods directly encode and concatenate the ambiguous Radar depth and images, confusing the learning pipeline and resulting in suboptimal depth estimation.

## III. METHODOLOGY

Our goal is to recover the dense depth  $\hat{\mathbf{d}} \in \mathbb{R}_+^{H_0 \times W_0}$  from a pair of RGB image  $\mathbf{I} \in \mathbb{R}^{3 \times H_0 \times W_0}$  and Radar point cloud  $\mathbf{P} = \{\mathbf{p}_i | \mathbf{p}_i \in \mathbb{R}^3, i = 0, 1, 2, \dots, k-1\}$  in

the image coordinate.  $H_0$  and  $W_0$  denote the height and width of the image, respectively. Either 3D or 4D Radar usually has a small field of view with ambiguous sparse data deteriorated by intensive noises. For cross-modal fusion, it is straightforward to project Radar points onto the image plane, generating Radar depth. However, direct fusion of the inherently ambiguous and sparse Radar depth with images, achieved by concatenating their encodings or raw data, can confuse the learning pipeline [12, 25–27], resulting in aliasing and other undesirable artifacts in the estimated depth. In this paper, we propose to get the scaleless dense depth with existing versatile monocular depth prediction networks, then learn to augment scaleless depth with accurate metric scales from Radar data.

The framework of our Radar-Camera depth estimation method consists of four stages: scaleless monocular depth prediction, global alignment (GA) of mono-depth, quasi-dense scale estimation, and scale map learner (SML) for refining dense scale locally, as shown in Fig.2.

### A. Monocular Depth Prediction

We employ off-the-shelf networks to predict robust and accurate scaleless depth from a single-view image. The high quality of the mono-depth prediction furnishes a solid foundation for scale-oriented learning. In this research, we harnessed SOTA mono-depth networks, like MiDaS v3.1 [3, 19] and DPT-Hybrid [20] with pre-trained weights on mixed diverse datasets. Both networks are built upon transformer architecture [28] and trained with scale and offset-invariant losses, ensuring strong generalization. They infer the relative depth relationship between pixels, producing dense depth (see Fig.3). Notably, our framework is versatile and compatible with arbitrary mono-depth prediction networks that predict depth  $\hat{\mathbf{d}}_m$ , inverse depth  $\hat{\mathbf{z}}_m$  or others.




Fig. 3: Left: the input image. Middle: Mono-Pred of MiDaS v3.1 [3]. Right: Mono-Pred of DPT-Hybrid [20]. Notably, MiDaS exhibits the ability to differentiate the sky.

### B. Global Alignment

We align the scaleless mono-depth prediction  $\hat{\mathbf{d}}_m$  with the Radar depth originating from projecting raw Radar points  $\mathbf{P}$ , by a global scaling factor  $\hat{s}_g$  and optional offset  $\hat{t}_g$ . The global aligned metric depth is calculated by  $\hat{\mathbf{d}}_{ga} = \hat{s}_g \cdot \hat{\mathbf{d}}_m + \hat{t}_g$ . Then, it is fed into the subsequent scale map learner (SML). There are many options for performing this global alignment between the projected Radar depth and mono-depth prediction, including: (i) **Var**: A varying  $\hat{s}_g$  for individual frame of mono-depth, calculated via root-finding algorithms [29, 30]. (ii) **Const**: A constant  $\hat{s}_g$  for all frames of mono-depth prediction, deemed as the mean of scale estimates on the entire training samples. (iii) **LS**:  $\hat{s}_g$  and  $\hat{t}_g$  for individual frames, computed with linear least-squares optimization [19]. (iv) **RANSAC**:  $\hat{s}_g$  and  $\hat{t}_g$  for

individual frames, computed with linear least-squares while incorporating RANSAC outlier rejection of the Radar depth. We randomly sample 5 Radar points with valid depth values, estimate  $\hat{s}_g$  and  $\hat{t}_g$  with the sampled Radar depth, and adopt the first pair of  $\hat{s}_g$  and  $\hat{t}_g$  that yields an inlier ratio over 90%. The inlier is the one where the discrepancy between the Radar point depth and aligned mono-depth is under 6m or the inverse depth discrepancy is under 0.015.

### C. Quasi-Dense Scale Estimation

Due to inherent sparsity and noises in Radar data, additional enhancement of raw Radar depth is crucial before conducting the scale map learner. To densify the sparse Radar depth obtained from projection, we exploit a transformer-based Radar-Camera data association network (shorthand RC-Net), which predicts the confidence of Radar-Pixel associations. Pixels without a direct correspondence of Radar point during projection might be associated with the depth of neighboring Radar point, thereby densifying the sparse Radar depth to a quasi-dense depth map, denoted as  $\hat{\mathbf{d}}_q$ .

1) *Network Architecture*: Our RC-Net (see Fig.2) is adapted from existing vanilla network RC-vNet [12] by further incorporating self and cross-attention [32] in a transformer module. The image encoder is a standard ResNet18 backbone [33] with 32, 64, 128, 128, 128 channels in each layer, and the Radar encoder is a multi-layer perceptron consisting of fully connected layers with 32, 64, 128, 128, 128 channels. The Radar features are mean pooled and reshaped to the shape of image features. Subsequently, Radar and image features are flattened and passed through  $N = 4$  layers of self and cross-attention, which involves a larger receptive field for the cross-modal association. These features, combined with skip connections from intermediate layers in the encoder, are forwarded to a decoder with logit output. Finally, the logits are activated by the sigmoid function to obtain the confidence map of cross-modal associations.

2) *Confidence of Cross-Modal Associations*: For a Radar point  $\mathbf{p}_i$  and a cropped image patch  $\mathbf{C}_i \in \mathbb{R}^{3 \times H \times W}$  in its projection vicinity, we use RC-Net  $h_\theta$  to obtain a confidence map  $\hat{\mathbf{y}}_i = h_\theta(\mathbf{C}_i, \mathbf{p}_i) \in [0, 1]^{H \times W}$ , which describes probability of whether the pixels in  $\mathbf{C}_i$  corresponds to  $\mathbf{p}_i$ . With  $k$  points in a Radar point cloud  $\mathbf{P}$ , the forward pass generates  $k$  confidence maps for individual Radar points. Therefore, each pixel  $\mathbf{x}_{uv}$ , ( $u \in [0, W_0 - 1], v \in [0, H_0 - 1]$ ) within image  $\mathbf{I}$  has  $n \in [0, k]$  associated Radar point candidates. By selecting the maximum score above the threshold  $\tau$ , we can find the corresponding Radar point  $\mathbf{p}_\mu$  for pixel  $\mathbf{x}_{uv}$ , and assign the depth of  $\mathbf{p}_\mu$  to  $\mathbf{x}_{uv}$ . Ultimately, this stage yields a quasi-dense depth map  $\hat{\mathbf{d}}_q \in \mathbb{R}_+^{H_0 \times W_0}$ :

$$\hat{\mathbf{d}}_q(u, v) = \begin{cases} d(\mathbf{p}_\mu), & \text{if } \hat{\mathbf{y}}_i(x_{uv}) > \tau \\ \text{None}, & \text{otherwise} \end{cases} \quad (1)$$

where  $\mu = \arg \max_i \hat{\mathbf{y}}_i(x_{uv})$ , and  $d(\cdot)$  returns the depth value. Finally, the quasi-dense scale map  $\hat{s}_q$  is calculated from  $\hat{s}_q = \hat{\mathbf{d}}_q / \hat{\mathbf{d}}_{ga}$ , and its inverse  $1/\hat{s}_q$  is subsequently fed into the scale map learner.

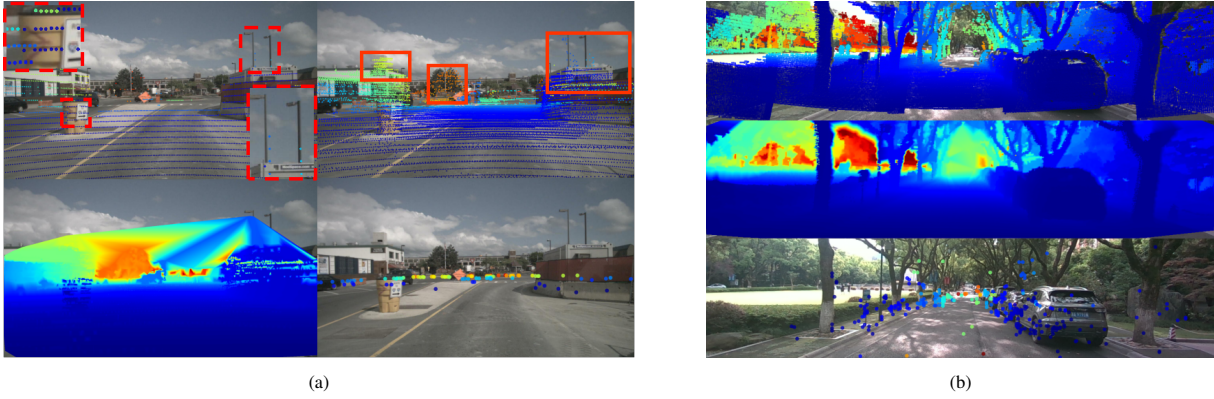


Fig. 4: (a) Top: nuScenes dataset [31] with LiDAR depth  $\mathbf{d}_{gt}$  and accumulated LiDAR depth  $\mathbf{d}_{acc}$ , depth from 3D Radar point cloud  $\mathbf{P}$ , interpolated LiDAR  $\mathbf{d}_{int}$  shown clockwise. The misalignment between LiDAR points and image pixels on this dataset is highlighted with red boxes. Depth from the 3D Radar point cloud is very sparse and non-uniformly distributed. (b) Our ZJU-4DRadarCam dataset with LiDAR depth  $\mathbf{d}_{gt}$ , interpolated LiDAR depth  $\mathbf{d}_{int}$ , and depth from 4D Radar point cloud  $\mathbf{P}$  shown from top to bottom. Compared to nuScenes, the ZJU-4DRadarCam dataset offers more accurate and denser LiDAR depth and denser 4D Radar depth.

3) *Training*: For the nuScenes dataset, we first project multiple frames to the current LiDAR frame  $\mathbf{d}_{gt}$  to obtain the cumulative LiDAR depth  $\mathbf{d}_{acc}$ . After that, linear interpolation in log space [34] is performed on  $\mathbf{d}_{acc}$  to obtain  $\mathbf{d}_{int}$ . Because of its density,  $\mathbf{d}_{gt}$  is directly interpolated without accumulation for the ZJU-4DRadarCam dataset. For supervision, we use  $\mathbf{d}_{int}$  to build binary classification labels  $\mathbf{y}_i \in \{0, 1\}^{H \times W}$ , where points with a depth difference less than 0.5m from the Radar point are labeled as positive. After constructing  $\mathbf{y}_i$ , we minimize the binary cross-entropy loss:

$$\mathcal{L}_{BCE} = \frac{1}{|\Omega|} \sum_{x \in \Omega} -(\mathbf{y}_i(x) \log \hat{\mathbf{y}}_i(x) + (1 - \mathbf{y}_i(x)) \log(1 - \hat{\mathbf{y}}_i(x))) \quad (2)$$

where  $\Omega \subset \mathbb{R}^2$  denotes the image region of  $\mathbf{C}_i$ ,  $x \in \Omega$  is a pixel coordinate, and  $\hat{\mathbf{y}}_i = h_\theta(\mathbf{C}_i, \mathbf{p}_i)$  is the confidence of correspondence.

#### D. Scale Map Learner

1) *Network Architecture*: Inspired by [5], we construct a scale map learner (SML) network based on MiDaS-small [19] architecture. SML aims to learn a pixel-level dense scaling map for  $\hat{\mathbf{d}}_{ga}$ , which completes the quasi-dense scale map and refine the metric accuracy of  $\hat{\mathbf{d}}_{ga}$ . SML requires concatenated  $\hat{\mathbf{z}}_{ga}$  and  $1/\hat{s}_q$  as input. The empty locations in  $\hat{s}_q$  are filled with ones. SML regresses a dense scale residual map  $\mathbf{r}$ , where values can be negative. We obtain the final scale map via  $1/\hat{\mathbf{s}} = \text{ReLU}(1 + \mathbf{r})$ , and the final metric depth estimation is computed by  $\hat{\mathbf{d}} = \hat{\mathbf{s}}/\hat{\mathbf{z}}_{ga}$ .

2) *Training*: Ground truth depth  $\mathbf{d}_{gt}$  is obtained from projecting 3D LiDAR points. LiDAR depth is further interpolated to get a densified depth  $\mathbf{d}_{int}$ . During training, we minimize the difference between the estimated metric dense depth  $\hat{\mathbf{d}}$  and  $\mathbf{d}_{gt}$ ,  $\mathbf{d}_{int}$  with a smoothed L1 penalty:

$$\mathcal{L}_{SML} = \mathcal{L}(\mathbf{d}_{int}, \hat{\mathbf{d}}) + \lambda_{gt} \mathcal{L}(\mathbf{d}_{gt}, \hat{\mathbf{d}}) \quad (3)$$

$$\mathcal{L}(\mathbf{d}, \hat{\mathbf{d}}) = \begin{cases} \frac{1}{|\Omega_d|} \sum_{x \in \Omega_d} (\mathbf{r}_d(x) - \beta/2), & \text{if } \mathbf{r}_d(x) < \beta \\ \frac{1}{|\Omega_d|} \sum_{x \in \Omega_d} (\mathbf{r}_d(x)^2 / 2\beta), & \text{otherwise} \end{cases} \quad (4)$$

where  $\mathbf{r}_d(x) = |\mathbf{d}(x) - \hat{\mathbf{d}}(x)|$ ,  $\lambda_{gt}$  is the weight of  $\mathcal{L}_{gt}$ ,  $\Omega_d \subset \Omega$  denotes the region where ground truth has valid depth values.  $\beta$  is set to 1 in our practice.

## IV. EXPERIMENTS

### A. Datasets

1) *NuScenes Dataset*: We first evaluate our method on nuScenes benchmark [31]. NuScenes dataset encompasses data collection across 1000 scenes in Boston and Singapore with LiDAR, 3D Radar, camera, and IMU sensors. It comprises around 40000 synchronized Radar-Camera keyframes. We followed the same data splits as [12] with 850 scenes for training and validation, and 150 for testing. The test split is officially offered by nuScenes v1.0.

2) *ZJU-4DRadarCam Dataset*: For extensive evaluation, we collected our own dataset, named ZJU-4DRadarCam, using a ground robot (Fig.6) equipped with Oculii's EAGLE 4D Radar, RealSense D455 camera, and RoboSense M1 LiDAR sensors. Our dataset consists of various driving scenarios, including urban and wilderness environments. Compared to nuScenes dataset, our ZJU-4DRadarCam offers 4D Radar data with denser measurements. Besides, we provide denser LiDAR depth for supervision and evaluation (see Fig.4). Our ZJU-4DRadarCam comprises a total of 33,409 synchronized Radar-Camera keyframes, split into 29312 frames for training and validation and 4097 frames for testing.

### B. Training Details and Evaluation Protocol

For the nuScenes dataset, following [12], we accumulate the individual  $\mathbf{d}_{gt}$  of 160 frames nearby to get  $\mathbf{d}_{acc}$ , which is then interpolated to yield  $\mathbf{d}_{int}$ . Dynamic objects are masked out during the above process. For the ZJU-4DRadarCam dataset, we directly interpolate  $\mathbf{d}_{gt}$  to obtain  $\mathbf{d}_{int}$  with linear interpolation [34] in the log space of depth.




Fig. 5: (a) Our metric depth estimation over the input image in a large-scale scenario. (b) Top row shows the ground truth depth  $d_{int}$  and Radar points  $P$  projected into image  $I$ . The rest rows from top to bottom depict the depth estimations of [11], [12], and our RadarCam-Depth and the corresponding error maps. Our method demonstrates much higher accuracy and fine details.



Fig. 6: Our data collection routes and the CAD model of our robot.

For training the RC-Net on nuScenes, with an input image size of  $900 \times 1600$ , the size of the cropped patch during confidence map formation is set to  $900 \times 288$ . For the training on ZJU-4DRadarCam, the input image size is  $300 \times 1280$ , while the patch size is  $300 \times 100$ . We employ the Adam optimizer with  $\beta_1=0.9$  and  $\beta_2=0.999$ , and a learning rate of  $2e^{-4}$  for 50 epochs. Data augmentations, including horizontal flipping, saturation, brightness, and contrast adjustments, are applied with a 0.5 probability. We train our RC-Net for 50 epochs with an NVIDIA RTX 3090 GPU, taking approximately 14 hours with a batch size of 6.

We adopt MiDaS-Small architecture for our SML, where the encoder backbone is initialized with pre-trained ImageNet weights [35], and other layers are randomly initialized. The input data is resized and cropped to  $288 \times 384$ . We use an Adam optimizer with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . The initial learning rate is set to  $2e^{-4}$  and reduced to  $5e^{-5}$  after 20 epochs. Training SML for 40 epochs takes about 24 hours with a batch size of 24.

Some widely adopted metrics from the literature are used for evaluating the depth estimations, including mean absolute error (MAE), root mean squared error (RMSE), absolute relative error (AbsRel), squared relative error (SqRel), the errors of inverse depth (iRMSE, iMAE), and  $\delta_1$  [36]. To better illustrate our experimental details, the demo video is available at <https://youtu.be/JDn0Sua5d9o>.

### C. Evaluation on NuScenes

We evaluate the metric dense depth against  $d_{gt}$  within the range of 50, 70, and 80 meters (see Tab.I). Our proposed RadarCam-Depth outperforms all the compared Radar-




Fig. 7: From left to right are the input images, depth estimation of [12], and depth estimation of our RadarCam-Depth on nuScenes benchmark. Ours accurately differentiates the sky region and preserves fine details on object boundaries.

Camera methods and surpasses the second best method [12] by a large margin at all ranges. Specifically, we observe 25.6%, 23.4%, and 22.5% reductions in MAE and 20.9%, 20.2%, and 19.4% drops in RMSE for 50m, 70m, and 80m, respectively. We attribute the outstanding performance of RadarCam-Depth to the reasonable monocular prediction  $\hat{d}_m$  and our scale learning strategy. Notably, RadarCam-Depth solely relies on a single-frame image and a Radar point cloud, obviating the need to aggregate multi-frame data. The ‘‘Radar’’ and ‘‘Image’’ columns in Tab.I specify the quantities of point clouds and images used as inputs for various methods. For nuScenes dataset, we adopt the sky-sensitive pre-trained model, MiDaS v3.1, as our monocular depth prediction network, which can accurately differentiate the sky from others (see Fig.3). However, since the depth estimations in sky regions are not associated with corresponding LiDAR depth, they are not counted during metric evaluations. Some snapshots of depth estimations from different methods are shown in Fig.7.

### D. Evaluation on ZJU-4DRadarCam

We follow a similar way to Sec.IV-C for the evaluations on the ZJU-4DRadarCam dataset. For the mono-depth prediction in our framework, we tried both MiDaS v3.1 [3], and DPT-Hybrid [20] models. The evaluations of the metric dense depth estimations from various methods are presented in Tab.II, where our approach at different configurations of mono-depth network and global alignment options are marked in bold. After a comprehensive evalu-

TABLE I: EVALUATIONS ON NUSCENES (mm)

Eval Dist	Method	Radar	Image	MAE	RMSE
50m	RC-PDA [10]	5	3	2225.0	4156.5
	RC-PDA with HG [10]	5	3	2315.7	4321.6
	DORN [11]	5(x3)	1	1926.6	4124.8
	Singh [12]	1	1	1727.7	3746.8
	<b>RadarCam-Depth</b>	1	1	<b>1286.1</b>	<b>2964.3</b>
70m	RC-PDA [10]	5	3	3326.1	6700.6
	RC-PDA with HG [10]	5	3	3485.6	7002.9
	DORN [11]	5(x3)	1	2380.6	5252.7
	Singh [12]	1	1	2073.2	4590.7
	<b>RadarCam-Depth</b>	1	1	<b>1587.9</b>	<b>3662.5</b>
80m	RC-PDA [10]	5	3	3713.6	7692.8
	RC-PDA with HG [10]	5	3	3884.3	8008.6
	DORN [11]	5(x3)	1	2467.7	5554.3
	Lin [8]	3	1	2371.0	5623.0
	R4Dyn [13]	4	1	N/A	6434.0
	Sparse-to-dense [37]	3	1	2374.0	5628.0
	PnP [38]	3	1	2496.0	5578.0
Singh [12]	1	1	2179.3	4898.7	
<b>RadarCam-Depth</b>	1	1	<b>1689.7</b>	<b>3948.0</b>	

ation, we observe that our methods with the DPT model perform better for depth metrics, while the methods with MiDaS demonstrate higher accuracy for inverse depth metrics. Overall, our proposed methodology exhibits significant improvements compared to existing Radar-Camera methods [12] and [11] (Fig.5). Compared to the second best [12], the best configuration of our method shows 40.2%, 40.1%, and 40.2% reductions in MAE within ranges of 50m, 70m, and 80m, respectively.

TABLE II: EVALUATIONS ON ZJU-4DRADARCAM (mm)

Dist	Method	MAE	RMSE	iMAE	iRMSE	AbsRel	SqRel	$\delta_1$
50m	DORN [11]	2210.171	4129.691	19.790	31.853	0.157	939.348	0.783
	Singh [12]	1785.391	3704.636	18.102	35.342	0.146	966.133	0.831
	DPT+Var+RC-vNet [12]	1243.339	3045.853	12.111	24.377	0.098	644.709	0.896
	<b>DPT+Const+RC-Net</b>	1082.927	<b>2803.180</b>	10.885	23.227	0.089	<b>561.834</b>	0.920
	<b>DPT+Var+RC-Net</b>	<b>1067.531</b>	2817.362	10.508	22.936	0.087	575.838	0.922
	MiDaS+Var+RC-Net	1177.257	3009.135	10.255	<b>22.385</b>	0.090	630.222	0.924
	MiDaS+LS+RC-Net	1083.691	2868.950	<b>10.059</b>	22.388	<b>0.086</b>	588.091	<b>0.928</b>
70m	DORN [11]	2402.180	4625.231	19.848	31.877	0.160	1021.805	0.777
	Singh [12]	1932.690	4137.143	17.991	35.166	0.147	1014.454	0.828
	DPT+Var+RC-vNet [12]	1337.649	3358.212	12.047	24.294	0.099	672.084	0.894
	<b>DPT+Const+RC-Net</b>	1178.046	3121.317	10.824	23.149	0.090	<b>589.377</b>	0.918
	<b>DPT+Var+RC-Net</b>	<b>1157.014</b>	<b>3117.721</b>	10.444	22.853	0.087	601.052	0.921
	MiDaS+Var+RC-Net	1280.124	3323.488	10.189	<b>22.300</b>	0.091	658.416	0.922
	MiDaS+LS+RC-Net	1177.253	3179.615	<b>9.996</b>	22.305	<b>0.086</b>	614.801	<b>0.926</b>
80m	DORN [11]	2447.571	4760.016	19.856	31.879	0.161	1038.919	0.776
	Singh [12]	1979.459	4309.314	17.971	35.133	0.147	1034.148	0.828
	DPT+Var+RC-vNet [12]	1365.383	3467.245	12.033	24.277	0.099	682.126	0.894
	<b>DPT+Const+RC-Net</b>	1206.541	3239.331	10.812	23.133	0.090	<b>599.674</b>	0.918
	<b>DPT+Var+RC-Net</b>	<b>1183.471</b>	<b>3229.999</b>	10.432	22.838	0.088	610.501	0.920
	MiDaS+Var+RC-Net	1309.859	3431.046	10.176	<b>22.282</b>	0.091	668.038	0.922
	MiDaS+LS+RC-Net	1205.137	3295.520	<b>9.984</b>	22.289	<b>0.086</b>	624.864	<b>0.926</b>

We report our proposed method’s runtime at the DPT-based mono-depth prediction configuration. The average processing times per frame are shown in Tab.III. Note that Mono-Pred and GA can run simultaneously with RC-Net. Regarding different scale global alignment methods, GA (Var) and GA (LS) exhibit relatively fast speeds, while GA (RANSAC) is significantly slow and not advocated.

TABLE III: RUNTIME TEST OF OUR MODULES (s)

Mono-Pred	GA (Const)	GA (Var)	GA (LS)	GA (RANSAC)	RC-Net	SML
0.0651	-	0.0624	0.0044	2.2903	0.2704	0.1227

### E. Ablation

1) *Transformer Module*: We commence our analysis by focusing on the transformer mechanism incorporated within our novel quasi-dense scale estimation module, RC-Net. When the transformer component is disengaged, the architecture is identical to the pre-existing vanilla network, RC-vNet

[12]. Our evaluation is conducted on the ZJU-4DRadarCam dataset, and the results are presented in Tab.IV. The comparison results are the error in quasi-dense depth estimation  $\hat{\mathbf{d}}_q$  against ground truth  $\mathbf{d}_{gt}$  within a range of 80 meters. Our RC-Net consistently outperforms RC-vNet across all evaluation metrics, as delineated in Tab.IV. Furthermore, as indicated in Tab.II, when integrated with the DPT+Var framework, DPT+Var+RC-Net exhibits notable performance superiority over its RC-vNet counterpart.

TABLE IV: ABLATION OF TRANSFORMER MODULE (mm)

Dataset	Method	MAE	RMSE	iMAE	iRMSE	Output Pts
ZJU-4D	RC-vNet [12]	1308.742	3339.697	20.418	38.540	172567.846
	RC-Net	<b>1083.305</b>	<b>3052.870</b>	<b>16.203</b>	<b>33.414</b>	<b>178713.881</b>

2) *Global Alignment*: Following the discussions in Sec.III-B, we systematically assess the four options for globally aligning the scale of mono-depth predictions. It is worth mentioning that we set a termination criterion of 400 iterations for the GA (RANSAC), at which point we halt the iterative process and select the values of  $\hat{s}_g$  and  $\hat{t}_g$  that yield the highest inlier ratio. The evaluation results of  $\hat{\mathbf{d}}_{ga}$  are presented in Tab.V within a range of 80 m. Combining the runtime performance in Tab.III, the best GA method for ZJU-4D (DPT) uses variable  $\hat{s}_g$  (Var), and the optimal method for ZJU-4D (MiDaS) is least-squares for  $\hat{s}_g$  and  $\hat{t}_g$  (LS). The experiments showcase that estimating  $\hat{t}_g$  for DPT leads to substantial inverse errors, significantly degrading the performance of the subsequent SML (conducted in inverse space). However, ZJU-4D (MiDaS) require simultaneous estimating  $\hat{s}_g$  and  $\hat{t}_g$  to achieve higher accuracy.

TABLE V: ABLATION OF GA MODULE (mm)

Data	Method	MAE	RMSE	iMAE	iRMSE	AbsRel	SqRel	$\delta_1$
ZJU-4D (DPT)	Const	4733.158	<b>6926.261</b>	37.942	53.913	0.392	<b>2501.383</b>	0.343
	Var	<b>4726.168</b>	6940.025	<b>36.531</b>	<b>52.569</b>	<b>0.386</b>	2569.835	<b>0.374</b>
	LS	5671.214	7409.278	111.292	530.436	0.552	4069.064	0.271
	RANSAC	5963.904	7662.980	336.568	1294.117	0.614	4732.633	0.277
ZJU-4D (MiDaS)	Const	14008.482	245011.138	38.973	51.691	0.752	25697452.430	0.358
	Var	7119.828	14297.549	32.285	46.340	0.468	13255.805	0.431
	LS	<b>4799.150</b>	<b>7968.478</b>	35.670	51.559	0.390	<b>4659.651</b>	0.394
	RANSAC	5113.080	11063.605	<b>23.920</b>	<b>37.881</b>	<b>0.347</b>	13322.258	<b>0.631</b>

## V. CONCLUSION

This paper presents a novel method for estimating dense metric depth by integrating monocular depth prediction with the scale from sparse and noisy Radar point clouds. We propose a dedicated four-stage framework that effectively combines the high-fidelity fine details of the image and the absolute scale of Radar data, surmounting the inherent challenges of detail loss and the imprecision of metrics that manifest in existing methods based on the direction fusion of Radar and image data or their encodings. Our experimental findings unequivocally demonstrate a significantly superior performance of the proposed methodology over the compared baseline, as substantiated by both quantitative and qualitative assessments. In general, we introduce a pioneering metric depth estimation solution, which is rigorously validated and suitable for application on fusing cameras with 3D or 4D Radars. In our future endeavors, we aim to enhance the applicability and effectiveness of our proposed method by leveraging vision foundation models pre-trained with abundant data.

## REFERENCES

- [1] Z. Teed and J. Deng. “Deepv2d: Video to depth with differentiable structure from motion”. In: *arXiv preprint arXiv:1812.04605* (2018).
- [2] A. Wong, X. Fei, S. Tsuei, and S. Soatto. “Unsupervised depth completion from visual inertial odometry”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 1899–1906.
- [3] R. Birkel, D. Wofk, and M. Müller. “MiDaS v3.1 – A Model Zoo for Robust Monocular Relative Depth Estimation”. In: *arXiv preprint arXiv:2307.14460* (2023).
- [4] V. Guizilini, I. Vasiljevic, D. Chen, R. Ambrus, and A. Gaidon. “Towards Zero-Shot Scale-Aware Monocular Depth Estimation”. In: *arXiv preprint arXiv:2306.17253* (2023).
- [5] D. Wofk, R. Ranftl, M. Müller, and V. Koltun. “Monocular Visual-Inertial Depth Estimation”. In: *arXiv preprint arXiv:2303.12134* (2023).
- [6] Y. Ma, X. Zhao, H. Li, Y. Gu, X. Lang, and Y. Liu. “RoLM: Radar on LiDAR Map Localization”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3976–3982.
- [7] Y. Ma, H. Li, X. Zhao, Y. Gu, X. Lang, L. Li, and Y. Liu. “FMCW Radar on LiDAR Map Localization in Structural Urban Environments”. In: (2023).
- [8] 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. 10233–10240.
- [9] Y. Long, D. Morris, X. Liu, M. Castro, P. Chakravarty, and P. Narayanan. “Full-velocity radar returns by radar-camera fusion”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 16198–16207.
- [10] 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. 12507–12516.
- [11] 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.
- [12] 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.
- [13] S. Gasperini, P. Koch, V. Dalletta, N. Navab, B. Busam, and F. Tombari. “R4Dyn: Exploring radar for self-supervised monocular depth estimation of dynamic scenes”. In: *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 751–760.
- [14] W. Zhao, S. Liu, Y. Shu, and Y.-J. Liu. “Towards better generalization: Joint depth-pose learning without posenet”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 9151–9161.
- [15] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia. “On the uncertainty of self-supervised monocular depth estimation”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 3227–3237.
- [16] L. Hoyer, D. Dai, Q. Wang, Y. Chen, and L. Van Gool. “Improving semi-supervised and domain-adaptive semantic segmentation with self-supervised depth estimation”. In: *International Journal of Computer Vision* (2023), pp. 1–27.
- [17] J. Bian, Z. Li, N. Wang, H. Zhan, C. Shen, M.-M. Cheng, and I. Reid. “Unsupervised scale-consistent depth and ego-motion learning from monocular video”. In: *Advances in neural information processing systems* 32 (2019).
- [18] X. Song, H. Hu, L. Liang, W. Shi, G. Xie, X. Lu, and X. Hei. “Unsupervised Monocular Estimation of Depth and Visual Odometry uUsing Attention and Depth-Pose Consistency Loss”. In: *IEEE Transactions on Multimedia* (2023).
- [19] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun. “Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer”. In: *IEEE transactions on pattern analysis and machine intelligence* 44.3 (2020), pp. 1623–1637.
- [20] R. Ranftl, A. Bochkovskiy, and V. Koltun. “Vision transformers for dense prediction”. In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2021, pp. 12179–12188.
- [21] K. Sartipi, T. Do, T. Ke, K. Vuong, and S. I. Roumeliotis. “Deep depth estimation from visual-inertial slam”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10038–10045.
- [22] X. Zuo, N. Merrill, W. Li, Y. Liu, M. Pollefeys, and G. Huang. “CodeVIO: Visual-inertial odometry with learned optimizable dense depth”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14382–14388.
- [23] J. Xie, C. Lei, Z. Li, L. E. Li, and Q. Chen. “Video depth estimation by fusing flow-to-depth proposals”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10100–10107.
- [24] 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.
- [25] Y. Yang, A. Wong, and S. Soatto. “Dense depth posterior (ddp) from single image and sparse range”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 3353–3362.
- [26] W. Van Gansbeke, D. Neven, B. De Brabandere, and L. Van Gool. “Sparse and noisy lidar completion with rgb guidance and uncertainty”. In: *2019 16th international conference on machine vision applications (MVA)*. IEEE, 2019, pp. 1–6.
- [27] A. Wong and S. Soatto. “Unsupervised depth completion with calibrated backprojection layers”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 12747–12756.
- [28] 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: *arXiv preprint arXiv:2010.11929* (2020).
- [29] G. E. Forsythe et al. *Computer methods for mathematical computations*. Prentice-hall, 1977.
- [30] R. P. Brent. *Algorithms for minimization without derivatives*. Courier Corporation, 2013.
- [31] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. “nusenes: A multimodal dataset for autonomous driving”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020, pp. 11621–11631.
- [32] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou. “LoFTR: Detector-free local feature matching with transformers”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2021, pp. 8922–8931.
- [33] 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.
- [34] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa. “The quick-hull algorithm for convex hulls”. In: *ACM Transactions on Mathematical Software (TOMS)* 22.4 (1996), pp. 469–483.

- [35] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. "Imagenet: A large-scale hierarchical image database". In: *2009 IEEE conference on computer vision and pattern recognition*. Ieee. 2009, pp. 248–255.
- [36] J. Sun, Y. Xie, L. Chen, X. Zhou, and H. Bao. "NeuralRecon: Real-time coherent 3D reconstruction from monocular video". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 15598–15607.
- [37] 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.
- [38] T.-H. Wang, F.-E. Wang, J.-T. Lin, Y.-H. Tsai, W.-C. Chiu, and M. Sun. "Plug-and-play: Improve depth estimation via sparse data propagation". In: *arXiv preprint arXiv:1812.08350* (2018).