

RaLiFlow: Scene Flow Estimation with 4D Radar and LiDAR Point Clouds

Jingyun Fu¹, Zhiyu Xiang^{1*}, Na Zhao^{2*}

¹Zhejiang University

²Singapore University of Technology and Design

fujingyun@zju.edu.cn, xiangzy@zju.edu.cn, na_zhao@sutd.edu.sg

Abstract

Recent multimodal fusion methods, integrating images with LiDAR point clouds, have shown promise in scene flow estimation. However, the fusion of 4D millimeter wave radar and LiDAR remains unexplored. Unlike LiDAR, radar is cheaper, more robust in various weather conditions and can detect point-wise velocity, making it a valuable complement to LiDAR. However, radar inputs pose challenges due to noise, low resolution, and sparsity. Moreover, there is currently no dataset that combines LiDAR and radar data specifically for scene flow estimation. To address this gap, we construct a Radar-LiDAR scene flow dataset based on a public real-world automotive dataset. We propose an effective preprocessing strategy for radar denoising and scene flow label generation, deriving more reliable flow ground truth for radar points out of the object boundaries. Additionally, we introduce RaLiFlow, the first joint scene flow learning framework for 4D radar and LiDAR, which achieves effective radar-LiDAR fusion through a novel Dynamic-aware Bidirectional Cross-modal Fusion (DBCF) module and a carefully designed set of loss functions. The DBCF module integrates dynamic cues from radar into the local cross-attention mechanism, enabling the propagation of contextual information across modalities. Meanwhile, the proposed loss functions mitigate the adverse effects of unreliable radar data during training and enhance the instance-level consistency in scene flow predictions from both modalities, particularly for dynamic foreground areas. Extensive experiments on the repurposed scene flow dataset demonstrate that our method outperforms existing LiDAR-based and radar-based single-modal methods by a significant margin.

Introduction

Scene flow estimation extracts 3D motion fields of the dynamic environment from consecutive input frames, ranging from monocular images (Liang et al. 2025), stereo pairs (Jiao, Tran, and Shi 2021), LiDAR point clouds (Kim et al. 2025) and emerging 4D radar data (Ding et al. 2022a). It provides class-agnostic information for scene understanding (Zhao, Chua, and Lee 2021; Li and Zhao 2024; Han et al. 2024; Wang et al. 2025) and supports higher-level tasks such as 3D detection (Sheng et al. 2022; Erçelik et al.

2022; Sheng et al. 2025), 3D panoptic segmentation (Pan et al. 2025), multi-object tracking (Najibi et al. 2022), autonomous navigation (Chen et al. 2025) and 4D occupancy forecasting (Guo et al. 2024). Recently, LiDAR-based scene flow estimation (Lin et al. 2024a; Luo et al. 2025) has become mainstream, with studies (Liu et al. 2022; Wan et al. 2023) integrating complementary information from 2D images or event cameras to enhance scene flow learning. However, since 2D data are not naturally aligned with 3D point clouds, these methods rely on 3D projection-based correspondences, leading to potential information loss during cross-modal fusion. Moreover, cameras and LiDAR sensors are susceptible to poor environmental conditions like storms and fog. In contrast, 4D millimeter wave radar, which shares the same point cloud format as LiDAR, is more robust to adverse weather due to its strong penetrative capabilities. Additionally, 4D radar is cost-effective and provides radial velocity measurement, offering direct motion cues for scene flow prediction. As such, 4D radar is a valuable complement to LiDAR for scene flow learning.

Despite the complementary nature of radar and LiDAR, their fusion remains unexplored for scene flow estimation due to two main challenges. First, radar-based scene flow learning is intractable due to the sparsity, inherent noise, and low resolution of the radar data (Fig. 1). Up to now, very few scene flow estimation methods (Ding et al. 2022a, 2023, 2024) have involved radar data, and its potential remains underutilized. Second, presently there is no real-world scene flow dataset that contains both radar and LiDAR data. Since manual scene flow labeling is quite expensive, some previous studies (Jund et al. 2021; Ding et al. 2022a, 2023; Li et al. 2023b; Zhang et al. 2024a) generate scene flow labels based on synchronized odometry data and rigid transformations between tracking boxes across consecutive frames. Although usable LiDAR scene flow labels can be produced from box-to-box correspondences, this processing method is suboptimal for radar because many dynamic radar points are outside the box, as shown in Fig. 2. These radar points will be incorrectly classified as static, resulting in unreliable flow labels and further negatively impacting scene flow learning in both training and evaluation stages.

To address the dataset gap, we derive a Radar-LiDAR scene flow dataset from VoD dataset (Palffy et al. 2022) with our effective preprocessing strategy. We propose a

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

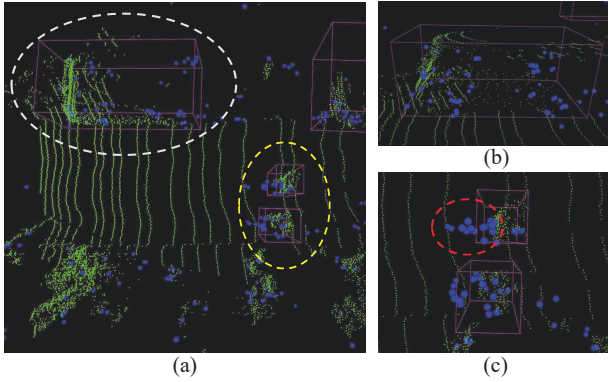


Figure 1: (a) Visualized comparison between 4D radar (blue) and LiDAR (green) point clouds. The magenta boxes are the tracked 3D bounding boxes from VoD dataset (Palffy et al. 2022) ground truth. (b) and (c) show the zoomed-in display results of the vehicle area (circled in white) and the pedestrian area (circled in yellow), respectively. It can be seen that the radar point cloud is much sparser and more noisy than LiDAR point cloud, with many radar points on the target object falling out of the bounding boxes (circled in red).

cross-modal hybrid radar denoising method that first identifies radar outliers by examining their absolute radial velocity (ARV). After that, the contextual information from the aligned LiDAR point cloud is used as a soft threshold to filter out isolated and irregular radar points. In addition, we propose a two-pronged strategy to generate more reliable radar scene flow labels. Specifically, we first calculate the rigid flow vectors for in-box points. Then, we identify the misclassified dynamic radar points based on the 3D spatial distance from each point to its nearest bounding box, the category label of the tracking box, and the consistency between ARV and the radial projection of the expected scene flow.

With our created dataset, it becomes possible to explore the radar-LiDAR fusion strategy for scene flow learning. Note that existing radar-LiDAR fusion methods are primarily designed for 3D object detection (Wang et al. 2022; Li et al. 2023a; Tai et al. 2024; Xu et al. 2024; Lin et al. 2024b) and cannot be directly applied to scene flow estimation. These methods focus on object-level feature extraction, while scene flow estimation requires fine-grained, class-agnostic dynamic details at the point level. Furthermore, the global attention mechanism (Liu, Shao, and Hoffmann 2021) commonly used in these methods tends to prioritize target object regions and often discard irrelevant scene details. This information loss can lead to significant performance degradation for scene flow prediction.

To address the method gap, we propose RaLiFlow, the first radar-LiDAR fusion framework for scene flow estimation, leveraging the strengths of both sensors. Pillar-based 2D features are extracted for 4D millimeter wave radar and LiDAR inputs, and complementary information interaction between these two modalities is enhanced with a novel Dynamic-aware Bidirectional Cross-modal Fusion (DBCF)

module. The DBCF module performs local cross-attention on the 2D feature maps of the two modalities under the guidance of radar dynamics. In particular, the velocity measurement inherent in radar data can initially indicate the dynamic foreground areas in the scene, which is most important and challenging for scene flow estimation. Inspired by this, the DBCF module computes a Gaussian heatmap based on the distance from each 2D grid center to its nearest dynamic radar grid, which provides attentive weight in the subsequent cross-attention operation. Moreover, we develop a set of loss functions to discard unreliable radar points in loss calculation and enforce scene flow coherence within dynamic foreground instances for both modalities.

Our main contributions are summarized as follows:

- We create a Radar-LiDAR scene flow dataset from VoD dataset with our effective preprocessing strategy, which provides more reliable radar scene flow labels.
- We present RaLiFlow, the first radar-LiDAR fusion network for scene flow estimation, exploiting the complementary information between the two modalities.
- We specially design a novel Dynamic-aware Bidirectional Cross-modal Fusion (DBCF) module which integrates radar dynamics into local cross-attention mechanism to enhance cross-modal information interaction.
- We develop effective loss functions to filter out unreliable radar supervision signals for training and encourage instance-level flow consistency of the two modalities.
- Extensive experiments on the repurposed scene flow dataset demonstrate that our method achieves state-of-the-art performance, surpassing existing single-modal methods by significant margins.

Related work

LiDAR Scene Flow Estimation. Nowadays, learning-based methods have been dominant for LiDAR scene flow estimation (Liu, Qi, and Guibas 2019; Wu et al. 2020; Cheng and Ko 2022, 2023; Wang et al. 2023a; Lin et al. 2024a) on synthetic datasets. However, these methods only support limited size of input data and suffer from the domain gap between synthetic and real data. To handle large-scale real-world LiDAR point clouds, a series of pillar-based (Jund et al. 2021; Zhang et al. 2024a,b; Khoche et al. 2025) and voxel-based (Li et al. 2022; Kim et al. 2025; Luo et al. 2025) pipelines rasterize world space into BEV grid and sparse volumetric grid, efficiently abstracting 2D and 3D grid-level features for the point clouds. Leveraging runtime optimization, another line of work (Li, Kaesemodel Pontes, and Lucey 2021; Li et al. 2023b; Vedder et al. 2023, 2024) is applicable to large-scale point cloud input but requires substantial computational resources. Moreover, several recent works (Liu et al. 2022; Wan et al. 2023) explore multimodal fusion for scene flow estimation and incorporate 2D modalities like images or event cameras with 3D LiDAR point clouds. However, how to fuse 4D millimeter wave radar and LiDAR data for scene flow estimation remains unknown.

Radar Scene Flow Estimation. Despite remarkable progress in LiDAR-based scene flow estimation, radar-

based scene flow estimation remains largely underexplored. RaFlow (Ding et al. 2022a) first introduces a self-supervised framework for scene flow prediction on radar point clouds, but suffers from the absence of effective supervision signals. MilliFlow (Ding et al. 2024) focuses on human motion estimation, relying on 3D human skeleton labels from RGB-D images for radar scene flow prediction. CMFlow (Ding et al. 2023) proposes a multitask architecture for radar scene flow learning, obtaining cross-modal supervision from ego motion, dynamic segmentation, optical flow and 3D scene flow labels generated by existing optical flow estimation and 3D tracking networks. Therefore, CMFlow is affected by the performance of these related models. Moreover, these existing radar-based methods can only handle a small number of input radar points. In contrast, our RaLiFlow can dynamically adapt to large scale real-world inputs and does not rely on existing networks for label generation.

3D Scene Flow Annotation. As annotating dense scene flow on real-world data is extremely expensive, synthetic datasets are widely used in early scene flow estimation methods (Geiger et al. 2013; Mayer et al. 2016; Jin et al. 2022). However, the distribution pattern of virtual point clouds differs from real data, and the commonly created point-level one-to-one correspondence between virtual frames is unrealistic. To generate flow ground truth from real-world datasets, previous methods (Jund et al. 2021; Ding et al. 2022b; Zhang et al. 2024a,b; Khoche et al. 2025; Li et al. 2023b) compute rigid scene flow from tracking boxes for in-box points and assign ego-motion vectors to points outside the boxes based on synchronized odometry. While this approach generates usable flow labels for LiDAR data, it is problematic for radar as many dynamic radar points may fall outside of the tracking boxes. As of now, though recent works (Lin and Caesar 2024; Jiang et al. 2024) have developed automatic flow annotation pipelines for LiDAR datasets, obtaining reliable scene flow ground truth for radar remains challenging.

Radar-LiDAR Fusion for 3D Detection. Recently, radar-LiDAR fusion (Wang et al. 2022, 2023b; Li et al. 2023a; Tai et al. 2024; Lin et al. 2024b) has proven effective in 3D object detection. However, these detection methods focus on object-level scene representation and the commonly used global attention may damage the point-level details in the scene, thus these methods cannot be simply applied to fine-grained scene flow estimation scenario. Currently, the dynamic cues inherent in radar have not been effectively exploited in cross-modal feature fusion.

Methodology

Given two consecutive frames of 4D radar and LiDAR data, let the radar point clouds in the source and target frames be denoted as \mathbf{P}^R and \mathbf{Q}^R , respectively, and the corresponding LiDAR point clouds as \mathbf{P}^L and \mathbf{Q}^L . The joint scene flow learning for both modalities aims to predict the point-wise scene flow vector $\mathbf{v}_i \in \mathbf{V}^R$ for each radar point $\mathbf{p}_i \in \mathbf{P}^R$, and $\mathbf{v}_j \in \mathbf{V}^L$ for each LiDAR point $\mathbf{p}_j \in \mathbf{P}^L$ in the source frame. The scene flow vectors can be decomposed into ego

vehicle’s motion and the absolute motion of each point:

$$\mathbf{V}^{R,L} = \mathbf{V}_{ego}^{R,L} + \hat{\mathbf{V}}^{R,L}. \quad (1)$$

For a fair comparison with previous studies (Jund et al. 2021; Zhang et al. 2024a; Li et al. 2023b), we assume that the vehicle’s ego motion is given and the output of our network is the relative scene flow $\hat{\mathbf{V}}^{R,L}$, which is later integrated with the known ego-motion flow to form the final scene flow prediction $\mathbf{V} = \{\mathbf{V}^R, \mathbf{V}^L\}$.

The goal of our task is to minimize the total End Point Error (EPE) between \mathbf{V} and the absolute ground truth flow \mathbf{V}^{gt} :

$$\min\left(\underbrace{\frac{1}{\|\mathbf{V}^R\|} \sum_{\mathbf{v}_i \in \mathbf{V}^R} \|\mathbf{v}_i - \mathbf{v}_i^{gt}\|_2}_{\text{Radar EPE}} + \underbrace{\frac{1}{\|\mathbf{V}^L\|} \sum_{\mathbf{v}_j \in \mathbf{V}^L} \|\mathbf{v}_j - \mathbf{v}_j^{gt}\|_2}_{\text{LiDAR EPE}}\right). \quad (2)$$

Radar-LiDAR-based Scene Flow Dataset

This section introduces how to create a Radar-LiDAR-based scene flow dataset from the real-world automotive VoD dataset (Palffy et al. 2022) with our effective data preprocessing strategy.

Ground Point Removal. Due to the aperture problem, estimating the local motion of points on flat ground surfaces is challenging. Therefore, ground removal is crucial for accurate scene flow estimation (Najibi et al. 2022; Lin and Caesar 2024; Zhang et al. 2024a,b; Khoche et al. 2025). In our approach, we first project the radar point clouds into the LiDAR coordinate system using the given sensor extrinsics. Then, we apply the method from (Himmelsbach, Hundelshausen, and Wuensche 2010) to remove the ground points from the combined radar and LiDAR scans.

Cross-Modal Hybrid Radar Denoising. After ground removal, we filter out radar outliers using a combination of a hard threshold based on the radar’s inherent motion information and spatial context (soft threshold) provided by the LiDAR data. The VoD dataset (Palffy et al. 2022) provides both relative and absolute radial velocity (ARV) measurements for each radar point. We establish a hard ARV threshold μ based on the average ARV of the dynamic radar points in each frame to filter out 1-10 abnormal radar points:

$$\mu = ARV_{avg} + \vartheta_{thre}. \quad (3)$$

where ϑ_{thre} is set to 10 m/s for computation efficiency and radar points with $ARV > 0.5$ m/s are identified as dynamic.

Additionally, we leverage contextual information from the aligned LiDAR data to remove isolated radar points, since radar point clouds are much sparser (500+ vs. 9000+ in our case) and noisier than LiDAR under normal conditions. Specifically, we project the combined point clouds onto a bird’s-eye view (BEV) grid with a grid size of 0.8 m. If no LiDAR points are found in a neighboring 3×3 grid area

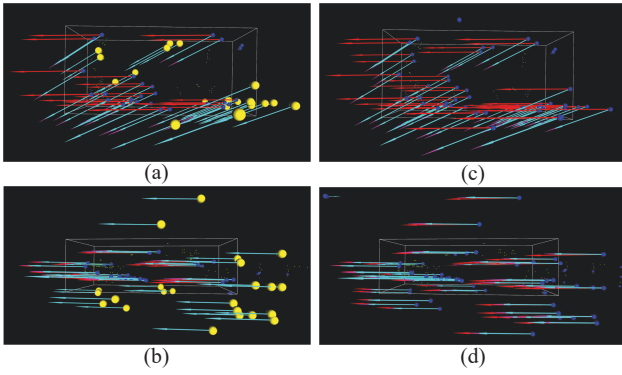


Figure 2: Comparison between the radar scene flow labels (red arrow) generated by our method (c and d) and previous methods (a and b). The illustration shows a moving car (white bounding box) from both bird’s-eye (a and c) and side views (b and d), with blue radar points and green LiDAR points correspond to it. The magenta arrows represent the radial projection of the generated radar scene flow labels, and cyan arrows represent point-wise ARV captured by the radar sensor. The highlighted yellow radar points indicate significant differences between the radial projection of absolute scene flow and ARV.

around a given radar point, that radar point is discarded. This loose search range ensures that only a few irregular radar points are removed and valid radar points are reserved even when the radar penetrates obstacles.

Scene Flow Label Generation. We generate in-box scene flow labels for both LiDAR and radar point clouds based on the rigid transformation between the corresponding tracking boxes in the frame sequence. In this process, points outside the labeled boxes are assigned flow values of 0 m/s. To address the noisy nature of radar points that may fall outside the boxes, we apply a two-pronged strategy to re-estimate the flow for those radar points outside the labeled boxes but with $ARV > 0.5$ m/s. First, we calculate the distance of each candidate radar point to the center of its nearest bounding box. Given that tracking boxes have different sizes across categories, we use the class labels in the VoD dataset to set category-specific distance thresholds based on statistical results from the training set. This allows us to determine whether a radar point is a valid dynamic point by considering its distance to the nearest bounding box and the class of the box. Additionally, we check the consistency between the ARV and the candidate radial flow for these radar points, with a conservative threshold of $\gamma_{thre} = 1$ m/s:

$$|radial(\mathbf{v}_i) - ARV_i| < \gamma_{thre}. \quad (4)$$

Candidate radar points that do not satisfy this condition are classified as static background points. As illustrated in Fig. 2, our two-pronged approach generates more reliable scene flow ground truth for radar inputs, mitigating the negative effects of poor radar label quality in both the training and evaluation phases.

RaLiFlow Network

Framework Architecture. Fig. 3 illustrates the overall framework of our proposed RaLiFlow. Given two consecutive point cloud frames of radar ($\mathbf{P}^R, \mathbf{Q}^R$) and LiDAR data ($\mathbf{P}^L, \mathbf{Q}^L$), we first extract a pseudo 2D feature map for each point cloud using pillar-based encoder (Desai, Schumann, and Alsheakhali 2020), denoted as $\phi(\cdot)$. A Gaussian heatmap generated from radar dynamics and a pair of pseudo 2D features from different modalities in each frame are then fed into our Dynamic-aware Bidirectional Cross-modal Fusion (DBCF) module. The DBCF module integrates radar dynamics into bidirectional local cross-attention between radar and LiDAR features to generate fused feature maps for each modality, denoted as $\varphi(\cdot)$. Here, $\varphi(\cdot)$ represents the combined operator of $\phi(\cdot)$ and the DBCF module. Subsequently, these fused features are then concatenated and fed into a U-Net backbone to produce a 2D flow embedding map. Finally, the flow embeddings, raw source point clouds, and concatenated fused feature maps from two frames are fed into two GRU-based flow heads (Cho et al. 2014; Zhang et al. 2024a), which iteratively generate the final scene flow predictions for each modality separately. Our RaLiFlow is trained using three loss functions: the LiDAR flow loss \mathcal{L}_{li} , the masked radar flow loss \mathcal{L}_{ra} , and the instance-level flow consistency loss \mathcal{L}_{ins} , as detailed in Sec. .

Dynamic-aware Bidirectional Cross-modal Fusion. To promote the interaction of complementary information across modalities, cross-attention provides a straightforward approach for fusing radar and LiDAR data. However, to better leverage the inherent motion information in 4D millimeter-wave radar and mitigate detail loss caused by the large receptive fields in global attention-based operations, we propose a novel Dynamic-aware Bidirectional Cross-modal Fusion (DBCF) module. This module integrates radar dynamics into a localized cross-attention mechanism. Specifically, radar dynamics-aware local cross-attention is applied bidirectionally: radar serves as the query to fuse local neighbors in the LiDAR domain, and vice versa, LiDAR serves as the query to fuse local neighbors in the radar domain.

Fig. 4 illustrates the detailed process of how to conduct the one-directional radar-to-LiDAR fusion. In this case, each non-sparse feature in the LiDAR feature map will be selected as query Q , and the non-sparse features in the 3×3 neighboring area of the radar feature map will be collected as keys K . After that, a relative 2D position encoding will be added to these neighboring features, yielding values V for the following cross-attention operation.

The absolute radial velocity (ARV) measurement from 4D radar naturally facilitates dynamic perception of the scene, allowing us to focus on challenging areas that are more likely to be dynamic when performing cross-attention operations. As shown in Fig. 4, the dynamic radar map represents pillars containing radar points with $ARV > 0.1$ m/s. We calculate the 2D distance D from each grid center to its nearest dynamic radar grid center, and the Gaussian heatmap G is then generated as follows:

$$G = \exp(-D^2/\sigma^2), \quad (5)$$

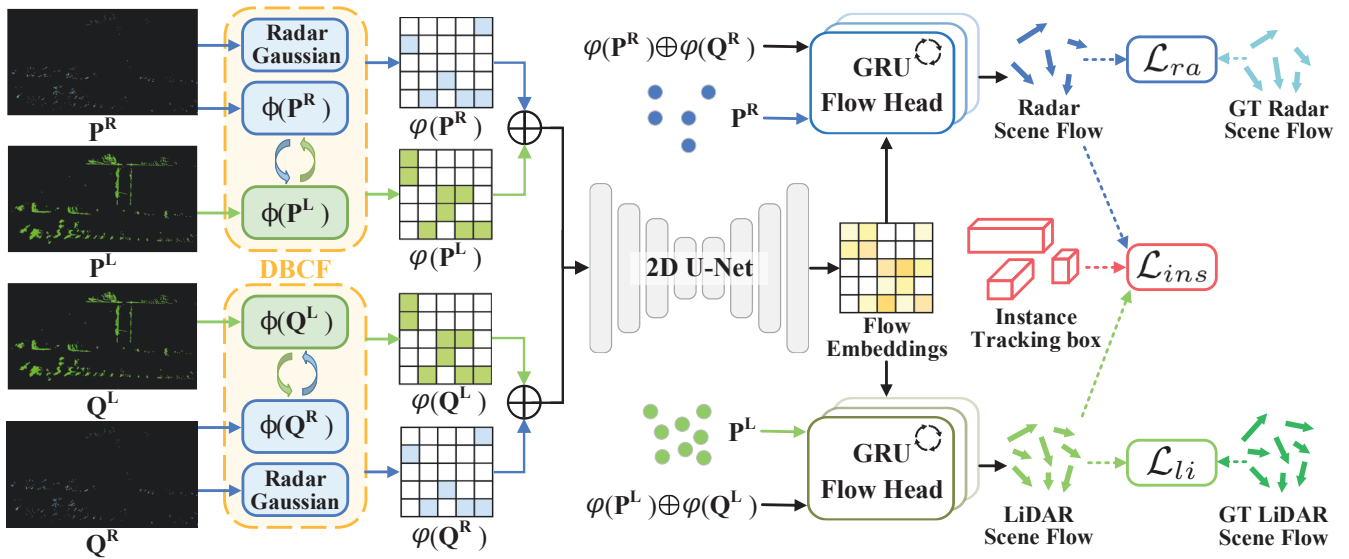


Figure 3: The overall architecture of RaLiFlow network. $\phi(\cdot)$ denotes pillar-based 2D feature map extracted for each input point cloud, and $\varphi(\cdot)$ represents the fused 2D feature map produced by DBCF module. GT stands for ground truth and \oplus represents concatenation.

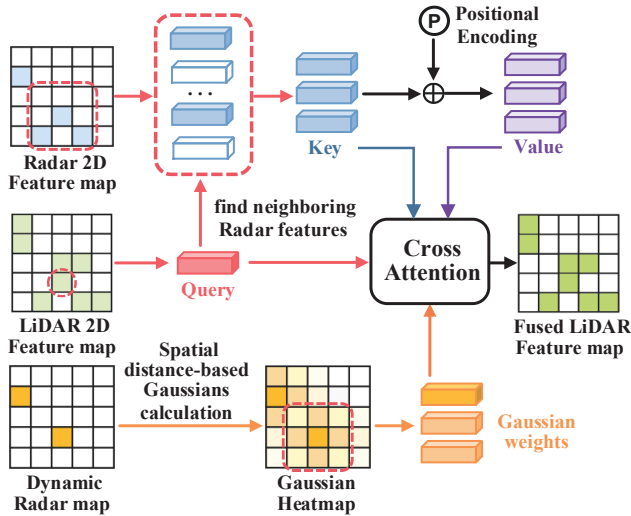


Figure 4: The process of one-directional dynamic-aware cross-modal fusion, with LiDAR feature serving as the query.

where σ determines the bandwidth of the Gaussian distribution. The generated G assigns higher weight to 2D features locate near the radar dynamic area when integrating modality \S into modality \dagger :

$$\mathcal{A}_{\S \rightarrow \dagger}(Q^\dagger, K^\S, V^\S | G) = \text{Softmax}\left(\frac{Q^\dagger (K^\S)^T}{\sqrt{c}} \odot G\right) V^\S, \quad (6)$$

where c is the channel dimension and \odot is the element-wise multiplication. Then the fused 2D feature maps $\varphi(\cdot)$ in each frame can be calculated from the previous pseudo 2D fea-

ture maps $\phi(\cdot)$. For the source frame, $\varphi(\mathbf{P}^R)$ and $\varphi(\mathbf{P}^L)$ are constructed as follows:

$$\varphi(\mathbf{P}^R) = \mathcal{A}_{L \rightarrow R}(Q^R, K^L, V^L | G) + \phi(\mathbf{P}^R), \quad (7)$$

$$\varphi(\mathbf{P}^L) = \mathcal{A}_{R \rightarrow L}(Q^L, K^R, V^R | G) + \phi(\mathbf{P}^L). \quad (8)$$

Loss Function

LiDAR Flow Loss. As most of the LiDAR points in real scenarios are static, we adopt the loss in (Zhang et al. 2024a) to balance the contribution of the input points. Since point-wise speed $\mathbf{s}(p_j)$ can be directly obtained from current flow prediction, the points in \mathbf{P}^L are divided into three point sets with different ranges of motion speeds:

$$p_j \in \begin{cases} \mathbf{S}_1^L & \text{if } \mathbf{s}(p_j) > 1.0 \text{ m/s} \\ \mathbf{S}_2^L & \text{if } \mathbf{s}(p_j) < 0.4 \text{ m/s} \\ \mathbf{S}_3^L & \text{o.w.} \end{cases} \quad (9)$$

where $\mathbf{S}_1^L \cup \mathbf{S}_2^L \cup \mathbf{S}_3^L = \mathbf{P}^L$. The LiDAR flow loss is calculated as:

$$\mathcal{L}_{li} = \sum_{k=1}^3 \frac{1}{|\mathbf{S}_k^L|} \sum_{p_j \in \mathbf{S}_k^L} \|\mathbf{V}^L(p_j) - \mathbf{V}_{gt}^L(p_j)\|_2. \quad (10)$$

Masked Radar Flow Loss. Since radar point clouds are much sparser and noisier, we filter out unreliable radar points with mismatched flow labels and ARV during loss calculation to reduce its negative impact on network training. Using Eq. (4), a point-wise confidence mask is generated for all radar points: $\mathbf{m}_i^r = \mathbf{1}(|\text{radial}(\mathbf{v}_{gt,i}) - \text{ARV}_i| < \gamma_{thre})$. $\mathbf{1}(\cdot)$ returns one when the condition in (\cdot) is satisfied. Similarly, the filtered radar points are divided into three

point sets: $\mathbf{S}_1^R \cup \mathbf{S}_2^R \cup \mathbf{S}_3^R = \{\mathbf{p}_i | \mathbf{m}_i^T = 1, \mathbf{p}_i \in \mathbf{P}^R\}$. Similar to Eq. (10), our masked radar flow loss is formulated as:

$$\mathcal{L}_{ra} = \sum_{k=1}^3 \frac{1}{|\mathbf{S}_k^R|} \sum_{\mathbf{p}_i \in \mathbf{S}_k^R} \|\mathbf{V}^R(p_i) - \mathbf{V}_{gt}^R(p_i)\|_2. \quad (11)$$

Instance-level Dynamic Flow Consistency Loss. Scene flow estimation on featureless surfaces is challenging because the network tends to match the points in the source frame to nearby areas with similar features in the target frame, which may lead to underestimated flow predictions. To address this issue, we suggest that the flow estimation results within each foreground instance should be consistent with the maximum flow among them. Specifically, we first identify H foreground instances $\{\mathbf{Ins}_h\}_{h=1}^H$ in the source frame based on the tracking boxes provided by the dataset. For each instance \mathbf{Ins}_h , we collect the corresponding radar points \mathbf{P}_h^R and LiDAR points \mathbf{P}_h^L located within it. Note that the valid dynamic radar points, as verified in Sec. , are also assigned to their nearest foreground instance. We then select and merge the dynamic points from both modalities to form a dynamic point set: $\mathbf{P}_h^D = \{p_c | \|\mathbf{V}_{gt}(p_c)\|_2 \geq 0.5 \text{ m/s}, p_c \in (\mathbf{P}_h^R \cup \mathbf{P}_h^L)\}$. Next, we find the index of the point in \mathbf{P}_h^D with the largest estimated scene flow:

$$\kappa = \arg \max_c \left\{ \|\mathbf{V}(p_c)\|_2 \mid p_c \in \mathbf{P}_h^D \right\}. \quad (12)$$

Based on this, we define the instance-level dynamic flow consistency loss as:

$$\mathcal{L}_{ins} = \sum_{h=1}^H \frac{1}{|\mathbf{P}_h^D|} \sum_{p_c \in \mathbf{P}_h^D} \|\mathbf{V}(p_c) - \mathbf{V}(p_\kappa)\|_2. \quad (13)$$

Finally, the total loss for RaLiFlow incorporates all three losses introduced above:

$$\mathcal{L}_{total} = \mathcal{L}_{li} + \mathcal{L}_{ra} + \mathcal{L}_{ins}. \quad (14)$$

Experiments

Experimental Setup

Datasets. With the data preprocessing procedure proposed in Sec. , we derive a Radar-LiDAR-based scene flow dataset from the View-of-Delft (VoD) dataset (Palffy et al. 2022). Currently, the VoD dataset is the only public automotive dataset that contains synchronized and calibrated real-world data of LiDAR and 4D radar point clouds, ego vehicle’s odometry, images, and 3D tracking box annotations frame by frame (total 8,600 frames). We follow the official split of the VoD dataset and generate scene flow labels for the training and validation set. The unlabeled test set is unused in our case. Following previous works (Ding et al. 2022a, 2023) on VoD dataset, only points within the viewing frustum are reserved, because the tracking annotation and radar data are not available in other regions.

Implementation Details. For a fair comparison with the previous works, we follow (Zhang et al. 2024a) and use four GRU iterations for flow prediction. We set $1/\sigma^2 = 10$

for computation efficiency. With a resolution of 0.1 m for each pillar encoder, we train our RaLiFlow and other pillar-based methods (Jund et al. 2021; Zhang et al. 2024a) for 150 epochs. To compare with voxel-based Flow4D (Kim et al. 2025), we implement both the official 2-frame version and a corresponding Flow4D* version, which sets the edge length of the voxel cube to 0.2 m and 0.1 m during voxelization, respectively. These voxel-based models are trained for 400 epochs. During training, we use Adam optimizer with an initial learning rate of 2×10^{-6} . In addition, we follow (Ding et al. 2023) to generate the extra required optical flow labels and train CMFlow on our proposed dataset under fully-supervised settings. Following official settings on VoD dataset, RaFlow (Ding et al. 2022a) and CMFlow are trained for 50 and 75 epochs, respectively. All the methods are trained on the VoD training dataset.

Evaluation Metrics. The End Point Error (EPE) measures the L2 norm of the discrepancy between the scene flow predictions and the ground truth flow vectors. Following previous works (Jund et al. 2021; Zhang et al. 2024a,b; Khoche et al. 2025; Kim et al. 2025), we adopt the three-way Endpoint Error (3-way EPE) for evaluation, which computes the unweighted average EPE of points located in Foreground Dynamic (\mathcal{FD}), Background Static (\mathcal{BS}) and Foreground Static (\mathcal{FS}) regions. For a convenient comparison with other existing methods which do not support 3-way EPE evaluation, we also present the commonly used 3D end-point-error (3D EPE) results, which is the averaged EPE result over all input points in the source frame.

Main Results

Quantitative Results. Since there is currently no other method that supports joint scene flow estimation on both radar and LiDAR point clouds, we compare our proposed method with other state-of-the-art single-modal methods in this section. Extensive comparisons are made with the pillar-based FastFlow3D (Jund et al. 2021) and DeFlow (Zhang et al. 2024a), the voxel-based Flow4D (Kim et al. 2025) as well as the point-based RaFlow (Ding et al. 2022a) and CMFlow (Ding et al. 2023). The experimental results in Table 1 demonstrate that our RaLiFlow achieves the leading position on VoD dataset with its novel and powerful radar-LiDAR fusion strategy and loss design. Compared with the radar-only baseline CMFlow (Ding et al. 2023), our RaLiFlow presents a significant performance improvement of 70.5% on the 3D EPE metric and a large error reduction of 50.2% on the 3-way EPE metric. Compared with LiDAR-only baselines, our RaLiFlow outperforms the best DeFlow (Zhang et al. 2024a) by a remarkable margin of 21.5% on the 3D EPE metric and 19.5% on the 3-way EPE metric. Note that our RaLiFlow can dynamically adapt to large scale real-world data inputs; RaFlow (Ding et al. 2022a) and CMFlow (Ding et al. 2023), on the other hand, only supports a fixed input size of 256 points. The quantitative results in Table 1 demonstrate the effectiveness of our RaLiFlow architecture and loss design.

Qualitative Results. In Fig. 5, we compare the qualitative outcomes of our RaLiFlow with scene flow groundtruth (GT) and the state-of-the-art single-modal methods. While

Method	Sup.	Input	LiDAR					Radar				
			Endpoint Error (\downarrow)					Endpoint Error (\downarrow)				
			3-way	3D	FD	BS	FS	3-way	3D	FD	BS	FS
FastFlow3D (Jund et al. 2021)	Full	L	0.0789	0.0491	0.1625	0.0348	0.0395	-	-	-	-	-
Flow4D (Kim et al. 2025)	Full	L	0.0771	0.0414	0.1708	0.0248	0.0359	-	-	-	-	-
Flow4D* (Kim et al. 2025)	Full	L	0.0732	0.0380	0.1650	0.0219	0.0326	-	-	-	-	-
DeFlow (Zhang et al. 2024a)	Full	L	0.0691	0.0321	0.1562	0.0152	0.0359	-	-	-	-	-
RaFlow (Ding et al. 2022a)	Self	R	-	-	-	-	-	0.1823	0.1678	0.2982	0.1458	0.1030
CMFlow (Ding et al. 2023)	Cross	R	-	-	-	-	-	0.1180	0.1242	0.1652	0.1187	0.0699
RaLiFlow (Ours)	Full	L+R	0.0556	0.0252	0.1282	0.0116	0.0269	0.0588	0.0366	0.1298	0.0206	0.0261

Table 1: Quantitative comparison on VoD validation set. RaFlow is self-supervised, and CMFlow requires cross-modal supervision from ego motion, dynamic segmentation, 2D optical flow and 3D scene flow labels; all the other methods are supervised by their input modalities. L denotes LiDAR input, and R denotes radar input.

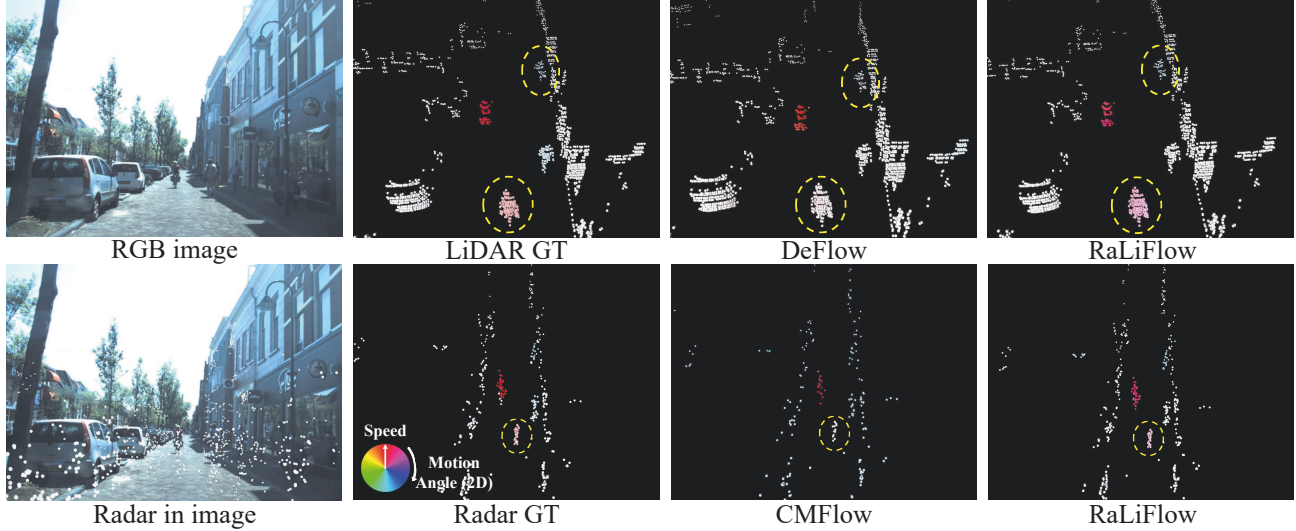


Figure 5: Qualitative Results on VoD validation dataset. The top row and the bottom row display the LiDAR-based and radar-based scene flow estimation results. The direction and magnitude of each flow vector are employed as hue and saturation, respectively. We enlarge the point size for better illustration, and the size of the radar point projected in the image reflects its depth in the 3D space.

DeFlow (Zhang et al. 2024a) fails to estimate the dynamic flows on pedestrians circled in yellow, our method generates more sensible scene flow prediction results. It can also be observed that CMFlow (Ding et al. 2023) is limited by the fixed size of the input radar point cloud and produces less reliable and sparser flows. Specifically, CMFlow (Ding et al. 2023) wrongly estimates the flows in the background and produces inaccurate flows for the pedestrian circled in yellow. Compared with other methods, our RaLiFlow generates scene flow predictions that best match the ground truth.

Ablation Studies

DBCF Module. We examine the design of DBCF module by replacing it with different radar-LiDAR fusion methods. As shown in Table 2, our DBCF module with bidirectional cross-attention operation outperforms the simple concatenation of radar and LiDAR 2D feature maps. Compared to simple concatenation, DBCF integrated with attentive Gaussian weights from radar dynamics achieves an overall performance improvement of 7.8% on the LiDAR 3-way EPE met-

ric and 7.3% on the radar 3-way EPE metric. Furthermore, the addition of radar dynamics improves the prediction accuracy of the DBCF in the most challenging \mathcal{FD} region as well as the \mathcal{BS} area, which illustrates that our novel DBCF better utilizes input radar data to promote a more comprehensive understanding of dynamic environments and achieve effective cross-modal fusion.

Loss Terms. The experimental results in Table 3 demonstrate that our masked \mathcal{L}_{ra} is able to reduce the adverse impact of unreliable radar points during training, enhancing the network performance in the difficult \mathcal{FD} region for radar. Moreover, incorporating the \mathcal{L}_{ins} constraint effectively reduces the scene flow estimation error on geometrically featureless moving surfaces, resulting in a significant performance improvement of 7.4% on the LiDAR 3D EPE metric and 8.7% on the radar 3D EPE metric.

Combination of Input Modalities. To further validate the effectiveness of our radar-LiDAR fusion method, we evaluate the radar-only and LiDAR-only performance under our

Fusion Method	Lidar					Radar				
	Endpoint Error (\downarrow)					Endpoint Error (\downarrow)				
	3-way	3D	FD	BS	FS	3-way	3D	FD	BS	FS
Concat	0.0603	0.0294	0.1358	0.0153	0.0297	0.0634	0.0408	0.1371	0.0242	0.0289
DBCF w/o Gaussian	0.0562	0.0261	0.1292	0.0127	0.0269	0.0596	0.0374	0.1312	0.0214	0.0262
DBCF with Gaussian	0.0556	0.0252	0.1282	0.0116	0.0269	0.0588	0.0366	0.1298	0.0206	0.0261

Table 2: Ablation study of Fusion methods on VoD validation set.

L_{ra}	L_{ins}	Lidar					Radar				
		Endpoint Error (\downarrow)					Endpoint Error (\downarrow)				
		3-way	3D	FD	BS	FS	3-way	3D	FD	BS	FS
masked		0.0577	0.0272	0.1314	0.0132	0.0286	0.0610	0.0401	0.1300	0.0243	0.0287
w/o mask	✓	0.0559	0.0254	0.1289	0.0118	0.0270	0.0596	0.0368	0.1320	0.0206	0.0263
masked	✓	0.0556	0.0252	0.1282	0.0116	0.0269	0.0588	0.0366	0.1298	0.0206	0.0261

Table 3: Ablation study of Loss functions on VoD validation set. \mathcal{L}_{ra} w/o mask simply sets $\mathbf{m}_i^r = 1$ for all radar points $\mathbf{p}_i \in \mathbf{P}^R$.

Input	Lidar					Radar				
	Endpoint Error (\downarrow)					Endpoint Error (\downarrow)				
	3-way	3D	FD	BS	FS	3-way	3D	FD	BS	FS
L+L	0.0673	0.0292	0.1555	0.0130	0.0333	-	-	-	-	-
R+R	-	-	-	-	-	0.0861	0.0561	0.1926	0.0325	0.0330
L+R	0.0562	0.0261	0.1292	0.0127	0.0269	0.0596	0.0374	0.1312	0.0214	0.0262

Table 4: Ablation study on different combination of input modalities on VoD validation set. ‘FD’ represents Foreground Dynamic, ‘BS’ represents Background Static, and ‘FS’ represents Foreground Static. L denotes LiDAR input, and R denotes radar input. RaLiFlow[†] is trained for 150 epochs with a batch size of 8 for this ablation.

proposed framework by duplicating one modality to preserve network structure. Since the generation of Gaussian heatmap requires radar information which is unavailable under LiDAR-only settings, we conduct this ablation study on RaLiFlow[†], which is another version of original RaLiFlow without Gaussian heatmap integrated in the DBCF module. For LiDAR-only input ablation, \mathcal{L}_{li} and \mathcal{L}_{ins} are employed for supervision; and for radar-only ablation, \mathcal{L}_{ra} and \mathcal{L}_{ins} are adopted for network training.

As shown in Table 4, the performance of our RaLiFlow[†] with LiDAR and radar input outperforms that of single-modal input on all evaluation metrics. Compared with LiDAR-only and radar-only cases, our radar-LiDAR fusion strategy improves the overall scene flow estimation performance on the 3-way EPE metric by 16.5% and 30.8%, respectively. This analysis verifies the complementary role of each sensor and the effectiveness of our cross-modal fusion design. Moreover, it can also be observed that the single-modal scene flow estimation results on our RaLiFlow[†] still outperform the corresponding prior state-of-the-art methods, demonstrating the advanced performance of our framework.

Conclusion

In this paper, we propose the first joint scene flow learning framework for 4D radar and LiDAR point clouds, namely RaLiFlow. We explore the complementary information interaction between these two modalities and achieve effective radar-LiDAR fusion via our novel DBCF module. We also

delve deeper into the inherent characteristics of 4D radar data and propose a novel data preprocessing strategy to generate more reliable radar scene flow labels based on the VoD dataset. The significant performance improvement of RaLiFlow over existing single-modal methods on the repurposed scene flow dataset demonstrates that our joint scene flow learning architecture effectively promotes the synergistic integration of 4D radar and LiDAR information. We hope our work highlights the promising potential of radar-LiDAR fusion for scene flow learning in the future.

Acknowledgements

This work was conducted during Jingyun’s visit to the IMPL Lab at SUTD, funded by China Scholarship Council. This research is supported by The Key Research & Development Plan of Zhejiang Province under Grant No.2024C01010, 2024C01017, 2025C01039, and the Joint R&D Program of the Yangtze River Delta Community of Sci-Tech Innovation with grant number 2024CSJGG01000. It is also supported by the Agency for Science, Technology and Research (A*STAR) under its MTC Programmatic Funds (Grant No. M23L7b0021).

References

Chen, Y.; Zhang, M.; Hao, Q.; and Zhou, G. 2025. SSF-PAN: Semantic Scene Flow-Based Perception for Au-

- onomous Navigation in Traffic Scenarios. *arXiv preprint arXiv:2501.16754*.
- Cheng, W.; and Ko, J. H. 2022. Bi-pointflownet: Bidirectional learning for point cloud based scene flow estimation. In *European Conference on Computer Vision*, 108–124. Springer.
- Cheng, W.; and Ko, J. H. 2023. Multi-scale bidirectional recurrent network with hybrid correlation for point cloud based scene flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10041–10050.
- Cho, K.; Van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Desai, N.; Schumann, T.; and Alsheikhali, M. 2020. A review of PointPillars architecture for object detection from point clouds. In *2020 IEEE International Conference on Consumer Electronics-Taiwan (ICCE-Taiwan)*, 1–2. IEEE.
- Ding, F.; Luo, Z.; Zhao, P.; and Lu, C. X. 2024. milliflow: Scene flow estimation on mmwave radar point cloud for human motion sensing. In *European Conference on Computer Vision*, 202–221. Springer.
- Ding, F.; Palffy, A.; Gavrilă, D. M.; and Lu, C. X. 2023. Hidden gems: 4d radar scene flow learning using cross-modal supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9340–9349.
- Ding, F.; Pan, Z.; Deng, Y.; Deng, J.; and Lu, C. X. 2022a. Self-supervised scene flow estimation with 4-d automotive radar. *IEEE Robotics and Automation Letters*, 7(3): 8233–8240.
- Ding, L.; Dong, S.; Xu, T.; Xu, X.; Wang, J.; and Li, J. 2022b. Fh-net: A fast hierarchical network for scene flow estimation on real-world point clouds. In *European Conference on Computer Vision*, 213–229. Springer.
- Erçelik, E.; Yurtsever, E.; Liu, M.; Yang, Z.; Zhang, H.; Topçam, P.; Listl, M.; Caylı, Y. K.; and Knoll, A. 2022. 3d object detection with a self-supervised lidar scene flow backbone. In *European Conference on Computer Vision*, 247–265. Springer.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237.
- Guo, E.; An, P.; Yang, Y.; Liu, Q.; and Liu, A.-A. 2024. FSF-Net: Enhance 4D Occupancy Forecasting with Coarse BEV Scene Flow for Autonomous Driving. *arXiv preprint arXiv:2409.15841*.
- Han, Y.; Zhao, N.; Chen, W.; Ma, K. T.; and Zhang, H. 2024. Dual-perspective knowledge enrichment for semi-supervised 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2049–2057.
- Himmelsbach, M.; Hundelshausen, F. V.; and Wuensche, H.-J. 2010. Fast segmentation of 3D point clouds for ground vehicles. In *2010 IEEE Intelligent Vehicles Symposium*, 560–565. IEEE.
- Jiang, C.; Wang, G.; Liu, J.; Wang, H.; Ma, Z.; Liu, Z.; Liang, Z.; Shan, Y.; and Du, D. 2024. 3dsflabelling: Boosting 3d scene flow estimation by pseudo auto-labelling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15173–15183.
- Jiao, Y.; Tran, T. D.; and Shi, G. 2021. Effiscene: Efficient per-pixel rigidity inference for unsupervised joint learning of optical flow, depth, camera pose and motion segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5538–5547.
- Jin, Z.; Lei, Y.; Akhtar, N.; Li, H.; and Hayat, M. 2022. Deformation and correspondence aware unsupervised synthetic-to-real scene flow estimation for point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7233–7243.
- Jund, P.; Sweeney, C.; Abdo, N.; Chen, Z.; and Shlens, J. 2021. Scalable scene flow from point clouds in the real world. *IEEE Robotics and Automation Letters*, 7(2): 1589–1596.
- Khoche, A.; Zhang, Q.; Sanchez, L. P.; Asefaw, A.; Mansouri, S. S.; and Jensfelt, P. 2025. SSF: Sparse Long-Range Scene Flow for Autonomous Driving. *arXiv preprint arXiv:2501.17821*.
- Kim, J.; Woo, J.; Shin, U.; Oh, J.; and Im, S. 2025. Flow4D: Leveraging 4D Voxel Network for LiDAR Scene Flow Estimation. *IEEE Robotics and Automation Letters*.
- Li, B.; Zheng, C.; Giancola, S.; and Ghanem, B. 2022. Sctn: Sparse convolution-transformer network for scene flow estimation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 1254–1262.
- Li, J.; Yang, L.; Chen, Y.; Yang, Y.; Jin, Y.; and Akiyama, K. 2023a. PillarDAN: Pillar-based Dual Attention Attention Network for 3D Object Detection with 4D RaDAR. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 1851–1857. IEEE.
- Li, L.; and Zhao, N. 2024. End-to-end semi-supervised 3d instance segmentation with pteacher. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 5352–5358. IEEE.
- Li, X.; Kaesemodel Pontes, J.; and Lucey, S. 2021. Neural scene flow prior. *NIPS*, 34: 7838–7851.
- Li, X.; Zheng, J.; Ferroni, F.; Pontes, J. K.; and Lucey, S. 2023b. Fast neural scene flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9878–9890.
- Liang, Y.; Badki, A.; Su, H.; Tompkin, J.; and Gallo, O. 2025. Zero-Shot Monocular Scene Flow Estimation in the Wild. *arXiv preprint arXiv:2501.10357*.
- Lin, M.; Xu, G.; Wang, Y.; Wang, X.; and Yang, X. 2024a. FlowMamba: Learning Point Cloud Scene Flow with Global Motion Propagation. *arXiv preprint arXiv:2412.17366*.
- Lin, Y.; and Caesar, H. 2024. Icp-flow: Lidar scene flow estimation with icp. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15501–15511.

- Lin, Z.; Liu, Z.; Xia, Z.; Wang, X.; Wang, Y.; Qi, S.; Dong, Y.; Dong, N.; Zhang, L.; and Zhu, C. 2024b. Rcbvdet: Radar-camera fusion in bird's eye view for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14928–14937.
- Liu, H.; Lu, T.; Xu, Y.; Liu, J.; Li, W.; and Chen, L. 2022. Camliflow: bidirectional camera-lidar fusion for joint optical flow and scene flow estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5791–5801.
- Liu, X.; Qi, C. R.; and Guibas, L. J. 2019. FlowNet3D: Learning scene flow in 3D point clouds. In *CVPR*, 529–537.
- Liu, Y.; Shao, Z.; and Hoffmann, N. 2021. Global attention mechanism: Retain information to enhance channel-spatial interactions. *arXiv preprint arXiv:2112.05561*.
- Luo, J.; Cheng, J.; Tang, X.; Zhang, Q.; Xue, B.; and Fan, R. 2025. MambaFlow: A Novel and Flow-guided State Space Model for Scene Flow Estimation. *arXiv preprint arXiv:2502.16907*.
- Mayer, N.; Ilg, E.; Hausser, P.; Fischer, P.; Cremers, D.; Dosovitskiy, A.; and Brox, T. 2016. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 4040–4048.
- Najibi, M.; Ji, J.; Zhou, Y.; Qi, C. R.; Yan, X.; Ettinger, S.; and Anguelov, D. 2022. Motion Inspired Unsupervised Perception and Prediction in Autonomous Driving. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVIII*, 424–443. Springer.
- Palfy, A.; Pool, E.; Baratam, S.; Kooij, J. F.; and Gavrilu, D. M. 2022. Multi-class Road User Detection with 3+ 1D Radar in the View-of-Delft Dataset. *RA-L*, 7(2): 4961–4968.
- Pan, Y.; Cui, Q.; Yang, X.; and Zhao, N. 2025. How Do Images Align and Complement LiDAR? Towards a Harmonized Multi-modal 3D Panoptic Segmentation. In *Forty-second International Conference on Machine Learning*.
- Sheng, H.; Cai, S.; Zhao, N.; Deng, B.; Huang, J.; Hua, X.-S.; Zhao, M.-J.; and Lee, G. H. 2022. Rethinking IoU-based optimization for single-stage 3D object detection. In *European Conference on Computer Vision*, 544–561. Springer.
- Sheng, H.; Cai, S.; Zhao, N.; Deng, B.; Liang, Q.; Zhao, M.-J.; and Ye, J. 2025. CT3D++: Improving 3D Object Detection with Keypoint-Induced Channel-wise Transformer. *International Journal of Computer Vision*, 133(7): 4817–4836.
- Tai, H.; Qian, Y.; Kang, X.; Liu, L.; and Liu, Y. 2024. Fusing LiDAR and Radar with Pillars Attention for 3D Object Detection. In *2024 7th International Symposium on Autonomous Systems (ISAS)*, 1–6. IEEE.
- Vedder, K.; Peri, N.; Chodosh, N.; Khatri, I.; Eaton, E.; Jayaraman, D.; Liu, Y.; Ramanan, D.; and Hays, J. 2023. ZeroFlow: scalable scene flow via distillation. *arXiv preprint arXiv:2305.10424*.
- Vedder, K.; Peri, N.; Khatri, I.; Li, S.; Eaton, E.; Kocamaz, M. K.; Wang, Y.; Yu, Z.; Ramanan, D.; and Pehserl, J. 2024. Scene Flow as a Partial Differential Equation. In *The Thirteenth International Conference on Learning Representations*.
- Wan, Z.; Mao, Y.; Zhang, J.; and Dai, Y. 2023. Rpeflow: Multimodal fusion of rgb-pointcloud-event for joint optical flow and scene flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10030–10040.
- Wang, L.; Zhang, X.; Xv, B.; Zhang, J.; Fu, R.; Wang, X.; Zhu, L.; Ren, H.; Lu, P.; Li, J.; et al. 2022. InterFusion: Interaction-based 4D radar and LiDAR fusion for 3D object detection. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 12247–12253. IEEE.
- Wang, X.; Zhao, N.; Han, Z.; Guo, D.; and Yang, X. 2025. Augrefer: Advancing 3d visual grounding via cross-modal augmentation and spatial relation-based referring. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 8006–8014.
- Wang, Y.; Chi, C.; Lin, M.; and Yang, X. 2023a. Ihnet: Iterative hierarchical network guided by high-resolution estimated information for scene flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10073–10082.
- Wang, Y.; Deng, J.; Li, Y.; Hu, J.; Liu, C.; Zhang, Y.; Ji, J.; Ouyang, W.; and Zhang, Y. 2023b. Bi-lrfusion: Bidirectional lidar-radar fusion for 3d dynamic object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13394–13403.
- Wu, W.; Wang, Z. Y.; Li, Z.; Liu, W.; and Fuxin, L. 2020. Pointpwc-net: Cost volume on point clouds for (self-) supervised scene flow estimation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, 88–107. Springer.
- Xu, R.; Xiang, Z.; Zhang, C.; Zhong, H.; Zhao, X.; Dang, R.; Xu, P.; Pu, T.; and Liu, E. 2024. SCKD: Semi-Supervised Cross-Modality Knowledge Distillation for 4D Radar Object Detection. *arXiv preprint arXiv:2412.14571*.
- Zhang, Q.; Yang, Y.; Fang, H.; Geng, R.; and Jensfelt, P. 2024a. DeFlow: decoder of scene flow network in autonomous driving. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2105–2111. IEEE.
- Zhang, Q.; Yang, Y.; Li, P.; Andersson, O.; and Jensfelt, P. 2024b. Seflow: A self-supervised scene flow method in autonomous driving. In *European Conference on Computer Vision*, 353–369. Springer.
- Zhao, N.; Chua, T.-S.; and Lee, G. H. 2021. Few-shot 3d point cloud semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8873–8882.