

CamLiFlow: Bidirectional Camera-LiDAR Fusion for Joint Optical Flow and Scene Flow Estimation

Haisong Liu* Tao Lu Yihui Xu Jia Liu Wenjie Li Lijun Chen
 State Key Laboratory for Novel Software Technology, Nanjing University, China
 {liuhs, taolu, yhxu, wenjielee}@smail.nju.edu.cn, {jialiu, chenlj}@nju.edu.cn

Abstract

In this paper, we study the problem of jointly estimating the optical flow and scene flow from synchronized 2D and 3D data. Previous methods either employ a complex pipeline that splits the joint task into independent stages, or fuse 2D and 3D information in an “early-fusion” or “late-fusion” manner. Such one-size-fits-all approaches suffer from a dilemma of failing to fully utilize the characteristic of each modality or to maximize the inter-modality complementarity. To address the problem, we propose a novel end-to-end framework, called CamLiFlow. It consists of 2D and 3D branches with multiple bidirectional connections between them in specific layers. Different from previous work, we apply a point-based 3D branch to better extract the geometric features and design a symmetric learnable operator to fuse dense image features and sparse point features. Experiments show that CamLiFlow achieves better performance with fewer parameters. Our method ranks 1st on the KITTI Scene Flow benchmark, outperforming the previous art with 1/7 parameters. Code is available at <https://github.com/MCG-NJU/CamLiFlow>.

1. Introduction

Optical flow and scene flow are the motion field in 2D and 3D space respectively. Through them, we can gain insights into the dynamics of the scene, which are critical to some high-level scene understanding tasks. In this work, we focus on the joint estimation of optical flow and scene flow, which addresses monocular camera frames with sparse depth measurements from LiDAR.

Previous methods [3, 31, 55, 56] construct a modular network that decomposes the estimation of flow into multiple subtasks. These submodules are independent of each other, making it impossible for utilizing their complementarity. Moreover, the limitations of any submodule will hurt the overall performance, since the whole pipeline depends on its results.

* Corresponding author (liuhs@smail.nju.edu.cn).

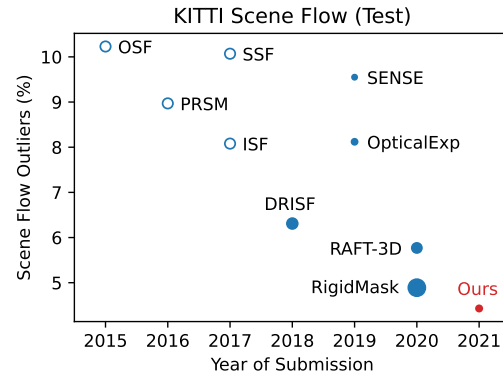


Figure 1. Results of the KITTI Scene Flow Benchmark. Marker size indicates model size. Models with unknown sizes and conventional approaches are marked as hollow. Our method outperforms all existing approaches [3, 23, 31, 33, 41, 45, 46, 55, 56] with much fewer parameters.

Other methods [42, 45] build an end-to-end architecture, consisting of multiple stages (including feature extraction, computing correlation, feature decoding, etc). RAFT-3D [45] concatenates images and dense depth maps to RGB-D frames and feeds them into a unified 2D network to predict pixel-wise 3D motion. This kind of “early fusion” (Fig. 2a) makes 2D CNNs fail to use most 3D structural information provided by depth. DeepLiDARFlow [42] takes images and LiDAR point clouds as input, where points are projected onto the image plane for densification and are fused with images in a “late-fusion” manner (Fig. 2b). However, some errors generated in the early stage lack the opportunity to be corrected by the other modality and are accumulated to the subsequent stages. Thus the complementarity between the two modalities is not fully exploited.

In general, single-stage fusion suffers from a dilemma of failing to fully utilize the characteristic of each modality or to maximize the inter-modality complementarity. To address the problem, we propose a **multi-stage** and **bidirectional** fusion pipeline (see Fig. 2c), which achieves better performance with fewer parameters. Within each stage, the two modalities are learned in separate branches using modality-specific architecture. At the end of each stage, a

learnable bidirectional bridge connects the two branches to pass complementary information. Moreover, recent point-based methods [28,29,36–38,48,51,52] achieve remarkable progress for 3D computer vision. This inspires us to process point clouds with a point-based branch, which can extract the fine 3D geometric information without any voxelization or projection.

It is worth noting that there are two challenges for the fusion of the image branch and the point branch. First, image features are organized in a dense grid structure, while point clouds do not conform to the regular grid and are sparsely distributed in the continuous domain. As a result, there is no guarantee of one-to-one correspondence between pixels and points. Second, LiDAR point cloud possesses the property of varying density, where nearby regions have much greater density than farther-away regions. To tackle the first problem, we propose a new learnable fusion operator, named *bidirectional camera-LiDAR fusion module* (Bi-CLFM), which fuses image/point features for both directions via learnable interpolation and sampling. As for the second problem, we propose a new transformation operator, named *inverse depth scaling* (IDS), which balances the distribution of points by scaling them non-linearly according to the inverse depth.

Experiments demonstrate that our approach achieves better performance with much fewer parameters. On FlyingThings3D [32], we achieve up to a **48.4%** reduction in end-point-error over RAFT-3D with only 1/6 parameters. On KITTI [33], CamLiFlow achieves an error of **4.43%**, outperforming the previous art [56] with only 1/7 parameters. The leaderboard is shown in Fig. 1.

2. Related Work

Optical Flow. Optical flow estimation aims to predict dense 2D motion for each pixel from a pair of frames. Traditional methods [4,6–8,16,50] often formulate optical flow as an energy minimization problem. FlowNet [12] is the first end-to-end trainable CNN for optical flow estimation, which adopts an encoder-decoder architecture. FlowNet2 [20] stacks several FlowNets into a larger one. PWC-Net [43] and some other methods [18,19,40,54] apply iterative refinement using coarse-to-fine pyramids. RAFT [44] constructs 4D cost volumes for all pairs of pixels and updates the flow iteratively. Although achieving state-of-the-art performance, RAFT runs much slower than PWC-Net. Hence, our two branches are built upon the PWC architecture to achieve a better trade-off between accuracy and speed.

Scene Flow from RGB-D Frames. RGB-D scene flow is the problem of estimating dense 3D motion for each pixel from a pair of stereo or RGB-D frames. Like optical flow, traditional methods [21,22,33,39] explore variational optimization and discrete optimization and treat

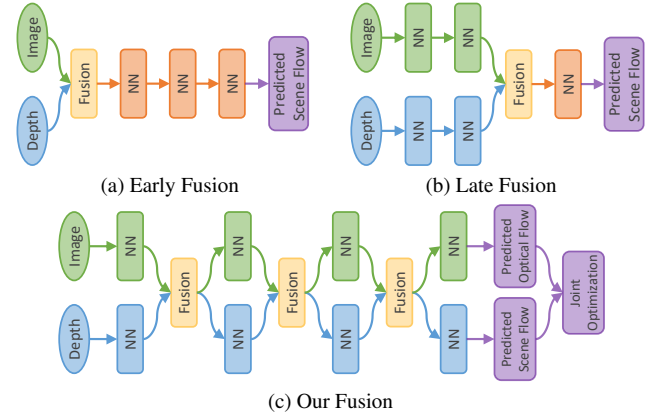


Figure 2. Architectures for feature-level fusion. Different from previous work which adopt an early/late fusion manner, we propose a multi-stage and bidirectional fusion pipeline.

scene flow as an energy minimization problem. Recent methods [3,31,55,56] divide scene flow estimation into multiple subtasks and build a modular network with one or more submodules for each subtask. Although achieving remarkable progress, their submodules are independent of each other, which can not exploit the complementary characteristics of different modalities. RAFT-3D [45] concatenate images and depth maps to RGB-D frames at an early stage, followed by a unified 2D network which iteratively updates a dense field of pixel-wise SE3 motion. However, this kind of “early fusion” makes it hard for 2D CNNs to take advantage of the rich 3D structural information.

Scene Flow from Point Clouds. Recently, researchers start to study scene flow estimation in 3D point clouds (e.g. from LiDAR) [14,25,28,29,35,48,49,52]. Based on [38], FlowNet3D [28] uses a flow embedding layer to represent the motion of points. FlowNet3D++ [48] achieves better performance by adding geometric constraints. Inspired by Bilateral Convolutional Layers, HPLFlowNet [14] projects the points onto a permutohedral lattice. PointPWC-Net [52] introduces a learnable cost volume for point clouds and estimates scene flow in a coarse-to-fine fashion. However, these methods do not exploit color features provided by images. As we demonstrate in our experiments, fusing point clouds with images can bring significant improvements.

Camera-LiDAR Fusion. Cameras and LiDARs have complementary characteristics, facilitating many computer vision tasks, such as depth estimation [13,30,57], scene flow estimation [2,42], 3D object detection [10,27,36,47,53], etc. Some researchers [2,36,47,57] build a modular network and perform result-level fusion, while the others [13,27,30,42,53] explore feature-level fusion schemes including early-fusion and late-fusion. Instead, we propose a multi-stage and bidirectional fusion pipeline, which not only fully utilizes the characteristic of each modality, but maximizes the inter-modality complementarity as well.

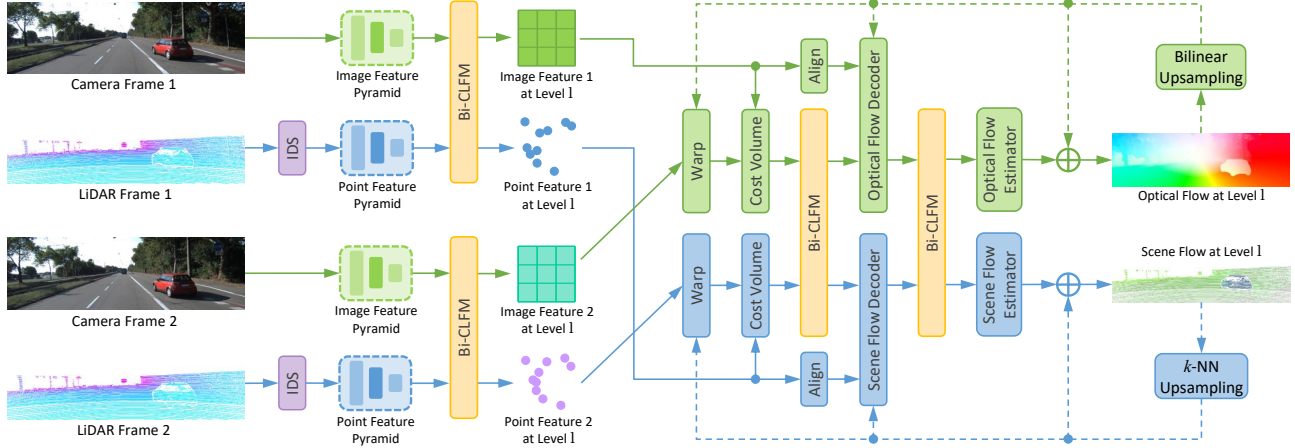


Figure 3. The architecture of CamLiFlow. Synchronized camera and LiDAR frames are taken as input, from which dense optical flow and sparse scene flow are estimated respectively. CamLiFlow is a two-branch network with multiple bidirectional fusion connections (Bi-CLFM) between them. Note that our method can also take advantage of other reliable depth sensors (such as stereo cameras).

3. CamLiFlow

Given a pair of the synchronized camera and LiDAR frames, CamLiFlow jointly estimates dense optical flow for camera frames and sparse scene flow for LiDAR frames. As illustrated in Fig. 3, CamLiFlow consists of two symmetric branches, named image branch and point branch, for 2D and 3D data respectively. Both branches are built on top of the PWC architecture [43, 52] where flow computed at the coarse level is upsampled and warped to a finer level. Features are fused in a bidirectional manner at multiple levels and stages.

In the following sections, we first introduce the *bidirectional camera-LiDAR fusion module* along with the multi-stage fusion pipeline. Next, we introduce *inverse depth scaling*, which makes the distribution of points more even across different regions. Finally, a multi-task loss for joint optimization is also introduced.

3.1. Bidirectional Camera-LiDAR Fusion Module

As mentioned above, the fusion between camera and LiDAR is challenging, since the data structures of image features and point features do not match. To overcome this, we introduce *bidirectional camera-LiDAR fusion module* (Bi-CLFM), which can fuse dense image features and sparse point features in a bidirectional manner.

As illustrated in Fig. 4, Bi-CLFM takes image features $F \in \mathbb{R}^{H \times W \times C_{2D}}$, point features $G = \{g_i | i = 1, \dots, N\} \in \mathbb{R}^{N \times C_{3D}}$ and point positions $P = \{p_i | i = 1, \dots, N\} \in \mathbb{R}^{N \times 3}$ as input, where N denotes the number of points. Features are fused for both directions so that both modalities can benefit each other. Note that we stop the gradient at specific locations to prevent one modality from dominating and stabilize the training (please refer to the supplementary material for more details).

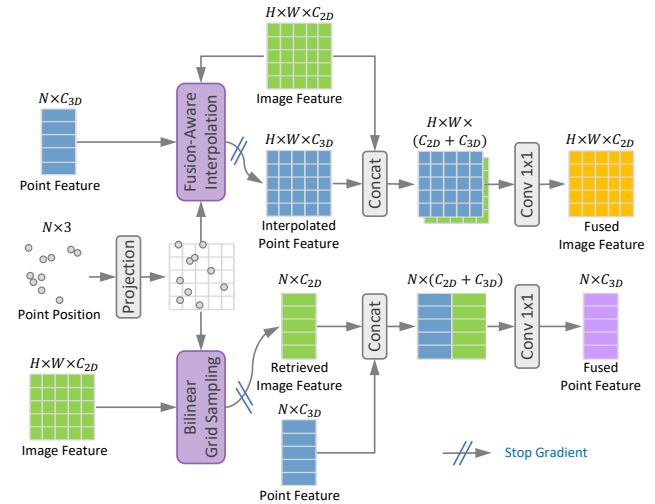


Figure 4. Details of Bidirectional Camera-LiDAR Fusion Module (Bi-CLFM). Features from two different modalities are fused in a bidirectional way, so that both modalities can benefit each other.

2D \Rightarrow 3D. First, points are projected to the image plane (denoted as $X = \{x_i | i = 1, \dots, N\} \in \mathbb{R}^{N \times 2}$) to retrieve the corresponding 2D feature:

$$H = \{F(x_i) | i = 1, \dots, N\} \in \mathbb{R}^{N \times C_{2D}}, \quad (1)$$

where $F(x)$ denotes the image feature at x and can be retrieved by bilinear interpolation if the coordinate is not a integer. Next, the retrieved feature H is concatenated with the input 3D feature G . Finally, a 1×1 convolution is employed to reduce the dimension of the fused 3D feature.

3D \Rightarrow 2D. Similarly, points are first projected to the image plane (denoted as $X = \{x_i | i = 1, \dots, N\} \in \mathbb{R}^{N \times 2}$). Since point clouds are sparse, we propose *fusion-aware interpolation* (detailed in the following paragraphs) to create a dense feature map $D \in \mathbb{R}^{H \times W \times C_{3D}}$ from sparse 3D fea-




Figure 5. Details of Fusion-Aware Interpolation. For each target pixel, we find the k nearest points around it. A learnable MLP followed by MEAN is employed to aggregate features.

tures. Next, the “interpolated” point features are concatenated with the input image features, followed by a 1×1 convolution to reduce the feature dimension.

Fusion-Aware Interpolation. To solve the problem of fusing sparse point features into dense image features, we propose a learnable *fusion-aware interpolation*. As illustrated in Fig. 5, for each target pixel q in the dense map, we find its k nearest neighbors among the projected points over the image plane. An MLP followed by MEAN is used to aggregate features, which can be formulated as:

$$D(q) = \frac{1}{k} \sum_{x_i \in \mathcal{N}_q} \text{MLP}([x_i - q, \mathcal{S}(q, x_i), g_i]), \quad (2)$$

where \mathcal{N}_q denotes all the neighborhood points, g_i is the 3D feature of point i and $[\cdot]$ denotes concatenation. The inputs of our MLP also include 2D similarity measurements between the q and its neighbors, which is defined as:

$$\mathcal{S}(q, x_i) = F(q) \cdot F(x_i). \quad (3)$$

Introducing 2D similarity measurements into the interpolation module makes it more robust in complex scenes with overlapping objects, since dense 2D features can be used to guide the densification of the sparse 3D features. We empirically test it in the ablation study (Fig. 9).

3.2. Multi-stage Fusion Pipeline

In this section, we build a multi-stage and bidirectional fusion pipeline with Bi-CLFM. Our backbone is based on the PWC architecture, which consists of multiple stages including feature extraction, warping, cost volume, and flow estimation. Within each stage, the two modalities are learned in separate branches using modality-specific architecture. At the end of each stage, a Bi-CLFM connects the two branches to pass complementary information.

Feature Pyramid. Given a pair of images and point clouds, we generate a feature pyramid for the image branch and the point branch respectively (the configuration details are included in the supplementary material). For each level l , image features are downsampled by a factor of 2 using residual blocks, while points are downsampled by the same

factor using *furthest point sampling*, followed by a PointConv [51] to aggregate features. The image pyramid encodes textural information, while the point pyramid encodes geometric information. Thus, features are fused by a Bi-CLFM at multiple levels for complementarity.

Warping. At each pyramid level l , both image features and point clouds are warped towards the reference frame using the upsampled flow from the lower level. Since the warping layer does not introduce any learnable parameters, we do not perform feature fusion after this stage.

Cost Volume. Cost volume stores the matching costs between the reference frame and the warped target frame. For the image branch, we follow [43] to construct a partial cost volume by limiting the search range to 4 pixels around each pixel. For the point branch, we follow [52] to construct a learnable cost volume layer. The pixel-based 2D cost volume maintains a fixed range of neighborhoods, while the point-based 3D cost volume searches for a dynamic range. Hence, we fuse the two cost volumes with a Bi-CLFM.

Flow Estimator. We build a flow estimator for each branch. The input of the flow estimator includes the cost volume, the features of the reference frame, and the upsampled flow. Our optical flow estimator follows [43], which employs a multi-layer CNN with DenseNet [17] connections. Our scene flow estimator follows [52], which is built as multiple layers of PointConv [51]. Features from the second last layer of the two estimators are fused. For clarity, we refer to the last layer as the “flow estimator” and the other layers as the “flow decoder” in Fig. 3.

3.3. Inverse Depth Scaling

As mentioned above, the distribution of LiDAR point clouds is not balanced, where nearby region has much greater density than farther-away region. Here, we propose a transformation operator for point clouds to address the problem, named *inverse depth scaling* (IDS). Formally, let (P_x, P_y, P_z) and (P'_x, P'_y, P'_z) be the coordinate of a point before and after the transformation respectively. IDS scales all three dimensions equally by the inverse depth $\frac{1}{P_z}$:

$$\frac{\delta P'_x}{\delta P_x} = \frac{\delta P'_y}{\delta P_y} = \frac{\delta P'_z}{\delta P_z} = \frac{1}{P_z}. \quad (4)$$

The transformed coordinates (P'_x, P'_y, P'_z) can be inferred by integrating the above formula:

$$P'_x = \int \frac{1}{P_z} dP_x = \frac{P_x}{P_z} + C_x, \quad (5)$$

$$P'_y = \int \frac{1}{P_z} dP_y = \frac{P_y}{P_z} + C_y, \quad (6)$$

$$P'_z = \int \frac{1}{P_z} dP_z = \log P_z + C_z, \quad (7)$$




Figure 6. The density of point clouds across different distances with/without IDS. We measure the local density around a point by averaging the offsets of its k nearest neighbors. IDS makes the distribution of points more even across different regions.

where both C_x and C_y are set to 0, and C_z is set to 1 to avoid zero depth.

In Fig. 6, we perform a statistic on FlyingThings3D and the raw Velodyne data of KITTI to show the density of points across different distances with/without IDS. The local density around a point is measured by averaging the offsets of its k nearest neighbors ($k = 16$). As we can see, IDS makes the distribution of points more even across different regions. In this paper, point clouds are transformed by IDS before being sent to the neural network.

3.4. Multi-task Loss

Although the estimation of optical flow and scene flow are highly relevant (the projection of scene flow onto the image plane becomes optical flow), we formulate them as two different tasks. We supervise the 2D and 3D branches respectively and design a multi-task loss for joint optimization. Let \hat{f}_{2D}^l and \hat{f}_{3D}^l be the ground truth optical flow and scene flow at the l th level respectively. The regression loss for each branch is defined as follows:

$$\mathcal{L}_{2D} = \sum_{l=l_0}^L \alpha_l \sum_{\mathbf{x}} \|f_{2D}^l(\mathbf{x}) - \hat{f}_{2D}^l(\mathbf{x})\|_2, \quad (8)$$

$$\mathcal{L}_{3D} = \sum_{l=l_0}^L \alpha_l \sum_{\mathbf{p}} \|f_{3D}^l(\mathbf{p}) - \hat{f}_{3D}^l(\mathbf{p})\|_2, \quad (9)$$

where $\|\cdot\|_2$ computes the L_2 norm. For fine-tuning, we use the following robust training loss:

$$\mathcal{L}_{2D} = \sum_{l=l_0}^L \alpha_l \sum_{\mathbf{x}} (|f_{2D}^l(\mathbf{x}) - \hat{f}_{2D}^l(\mathbf{x})| + \epsilon)^q, \quad (10)$$

$$\mathcal{L}_{3D} = \sum_{l=l_0}^L \alpha_l \sum_{\mathbf{x}} (|f_{3D}^l(\mathbf{x}) - \hat{f}_{3D}^l(\mathbf{x})| + \epsilon)^q, \quad (11)$$

where $|\cdot|$ computes the L_1 norm, $q = 0.4$ gives less penalty to outliers and ϵ is set to 0.01. The final loss is a weighted sum of the losses defined above:

$$\mathcal{L} = \mathcal{L}_{2D} + \lambda \mathcal{L}_{3D}, \quad (12)$$

where λ is set to 1.0 for all our experiments.

4. Experiments

We implement our model using PyTorch [34]. For all experiments we use the Adam optimizer [24] with weight decay set to 10^{-6} . The loss weights are set to $\alpha_0 = 8$, $\alpha_1 = 4$, $\alpha_2 = 2$, $\alpha_3 = 1$, and $\alpha_4 = 0.5$.

4.1. Main Results

We evaluate our method on the synthetic dataset FlyingThings3D [32] and the real-world dataset KITTI [33]. FlyingThings3D consists of stereo and RGB-D images rendered with multiple randomly moving objects from ShapeNet [9], which is large-scale and challenging. KITTI Scene Flow is a real-world benchmark for autonomous driving, consisting of 200 training scenes and 200 test scenes.

4.1.1 FlyingThings3D

Data Preprocessing. Following previous work [14, 20, 52], we use the subset of FlyingThings3D. The training and validation set respectively contains 19640 and 3824 pairs of camera-LiDAR frames. We follow FlowNet3D [28] instead of HPLFlowNet [14] to lift the depth images to point clouds, since HPLFlowNet only keeps non-occluded points which oversimplifies the problem.

Training. The training consists of two stages. First, we train our model for 600 epochs with the L_2 -norm loss function. The initial learning rate is set to 4×10^{-4} and reduced by half at 400 and 500 epochs. Next, we fine-tune our model for another 800 epochs with the robust loss function and a fixed learning rate of 10^{-4} . The batch size is set to 32.

Evaluation Metrics. Following RAFT-3D, we evaluate our network using 2D and 3D end-point error (EPE), as well as threshold metrics ($\text{ACC}_{1\text{px}}$ and $\text{ACC}_{0.05}$), which measure the portion of error within a threshold.

Quantitative Results. In Tab. 1, we compare to several state-of-the-art methods which utilize different input modalities. By fusing the two modalities of camera and LiDAR, our method outperforms all image-only and LiDAR-only methods by a large margin. Our method also outperforms RAFT-3D, which has 45M parameters and takes dense RGB-D frames as input. In contrast, our model is much more lightweight with 7.7M parameters and only requires sparse depth measurements. Moreover, our model reduces the best published EPE_{3D} from 0.062 to 0.032, which proves the superior performance of the point branch.

Qualitative Results. The visual comparison of optical flow and scene flow estimation is shown in Fig. 7. We also add two single-modal variations of our method for comparison, which removes the 2D branch or the 3D branch. As we can see, our full model better handles objects with repetitive structures and complex scenes with overlapping objects.

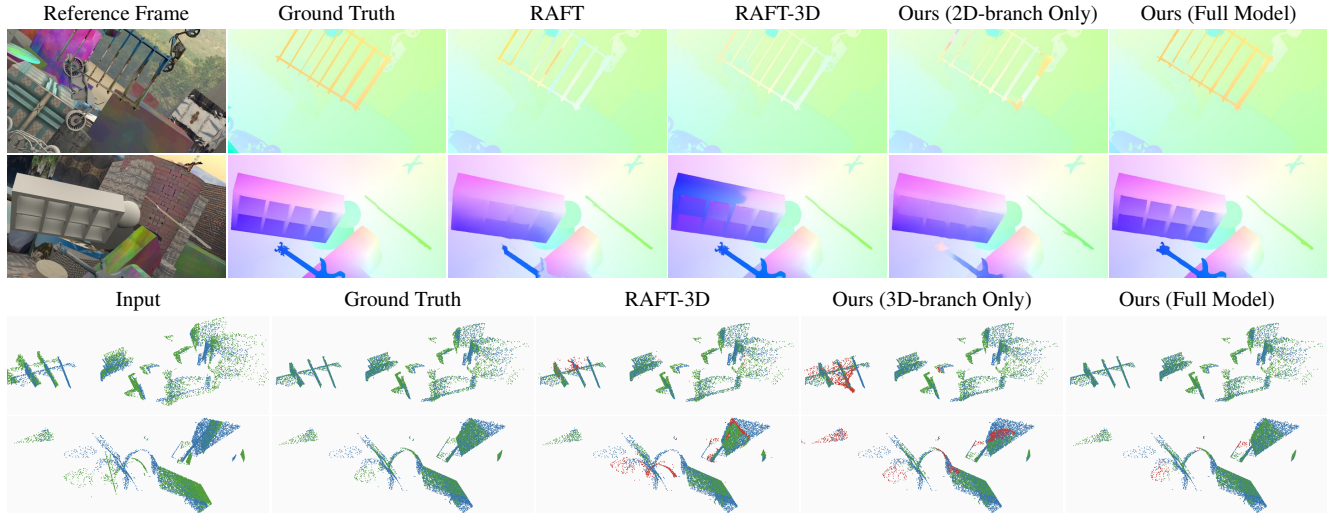


Figure 7. Visualized optical flow and scene flow estimation on the “val” split of the FlyingThings3D subset. The outliers are marked as red for scene flow estimation. Our full model better handles objects with repetitive structures and texture-less regions.

Method	Input	2D Metrics		3D Metrics		Parameters
		EPE _{2D}	ACC _{1px}	EPE _{3D}	ACC _{.05}	
FlowNet2.0 [20]	Image	5.05	72.8%	-	-	162.5M
PWC-Net [43]	Image	6.55	64.3%	-	-	9.4M
RAFT [44]	Image	3.12	81.1%	-	-	5.3M
FlowNet3D [28]	LiDAR	-	-	0.151	20.7%	1.2M
PointPWC [52]	LiDAR	-	-	0.112	51.8%	5.3M
FLOT [35]	LiDAR	-	-	0.170	23.4%	0.1M
DeepLiDARFlow [42]	Image+LiDAR	6.04	47.1%	-	27.2%	8.3M
RAFT-3D [45]	Image+Depth	2.37	87.1%	0.062	84.5%	45M
Ours (W/o fine-tuning)	Image+LiDAR	2.18	84.3%	0.033	91.4%	7.7M
Ours	Image+LiDAR	2.20	87.3%	0.032	92.6%	7.7M

Table 1. Performance comparison on the “val” split of the FlyingThings3D subset. For 2D metrics, we evaluate on full images excluding extremely fast moving regions with flow > 250 px. For 3D metrics, we follow the setup of FlowNet3D in which only non-occluded points with depth < 35 m are considered for evaluation.

4.1.2 KITTI

Training. Using the weight pre-trained on FlyingThings3D and Driving [32], we fine-tune our model on KITTI for 300 epochs with a fixed learning rate of 5×10^{-5} and a batch size of 8. We follow [55, 56] and divide the 200 training images into *train*, *val* splits based on the 4:1 ratio. During training, we lift the ground-truth disparity maps into point clouds using the provided calibration parameters. Basic data augmentation strategies including color jitter, random horizontal flipping, and random cropping are applied.

Testing. During testing, since neither disparity maps nor point clouds are provided, we employ GA-Net [58] to estimate the disparity from stereo images, and generate point clouds with depth < 90 m. The sparse output of our point branch is interpolated to create a dense prediction.

Refinement of Background Scene Flow. Since most background objects in KITTI are rigid (e.g. ground, build-

ings, etc), we can refine the background scene flow using a rigidity refinement step. Specifically, we employ DDRNet-Slim [15], a light-weight 2D semantic segmentation network, to determine the rigid background. DDRNet-Slim is pre-trained on Cityscapes [11] and fine-tuned on KITTI. Next, we estimate ego-motion by fitting and decomposing essential matrices from the background flow map using a neural-guided RANSAC [5]. Finally, the background scene flow is refined using the ego-motion and the disparity of the first frame.

Comparison with State-of-the-art Methods. We submit our approach to the website of KITTI Scene Flow benchmark and report the leaderboard in Tab. 2. A visualized comparison is shown in Fig. 8. Our approach outperforms all published methods, including RigidMask [56] (SF-all: 4.43% vs. 4.89%), which employs more than 140M parameters. In contrast, our method is much more lightweight with only 19.7M parameters (6.3M GA-Net + 7.7M Cam-

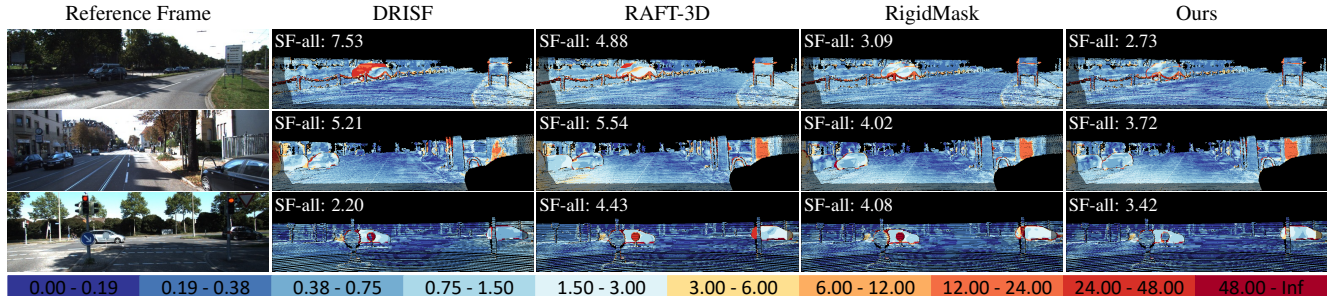


Figure 8. Qualitative results on the KITTI Scene Flow test set. Blue indicates a lower error, red indicates a higher error. Our approach improves accuracy near motion boundaries.

Method	Rigidity	D1 (%)		D2 (%)		OF (%)		SF (%)		Parameters
		noc	all	noc	all	noc	all	noc	all	
SENSE [23]	Background	2.05	2.22	3.87	5.89	5.98	7.64	7.30	9.55	13.4M
Binary TTC [1]	None	1.63	1.81	2.72	4.76	3.89	6.31	5.29	8.50	-
OpticalExp [55]	None	1.63	1.81	2.62	4.25	3.89	6.30	5.21	8.12	18.5M
ISF [3]	Full Image	4.02	4.46	4.95	5.95	4.69	6.22	6.45	8.08	-
ACOSF [26]	Full Image	3.35	3.58	4.26	5.31	4.51	5.79	6.40	7.90	-
DRISF [31]	Full Image	2.35	2.55	3.14	4.04	3.58	4.73	4.99	6.31	58.9M
RAFT-3D [45]	Full Image	1.63	1.81	2.67	3.67	3.23	4.29	4.53	5.77	51.3M
RigidMask [56]	Full Image	1.70	1.89	2.47	3.23	2.54	3.50	3.73	4.89	145.3M
Ours (non-rigid)	None	1.63	1.81	2.39	3.19	2.77	4.05	4.03	5.62	14.0M
Ours	Background	1.63	1.81	2.37	2.95	2.40	3.10	3.55	4.43	19.7M

Table 2. Leaderboard of the KITTI Scene Flow benchmark. “D1”, “D2”, and “OF” are predicted by GA-Net, our point branch and our image branch respectively. Our method ranks first on the leaderboard and can handle general non-rigid motions since we only treat the static background as rigid.

LiFlow + 5.7M DDRNet-Slim). Moreover, previous methods leverage more strict rigid-body assumptions by assigning rigid motions to all objects, while our method can handle general non-rigid motions since we only apply rigid motion refinement to the static background.

If the rigidity refinement step of the background scene flow is removed (corresponding to our “non-rigid” variation in Tab. 2), our method still ranks second on the leaderboard (SF-all: 5.62%). In this setting, our method does not require the background segmentation labels and can deal with any non-rigid motion (no matter foreground or background). Instead, RigidMask fails to handle non-rigid motions and suffers from the limitations of the motion segmentation network, since the whole pipeline depends on its results.

4.2. Ablation Study

In this section, we conduct ablation studies on FlyingThings3D to confirm the effectiveness of each module. All variations are trained for the first stage without fine-tuning with the robust loss function.

Unidirectional Fusion vs. Bidirectional Fusion. CamLiFlow fuses features in a bidirectional manner. Here, we train two variations where features are fused in a unidirectional manner ($2D \Rightarrow 3D$ or $2D \Leftarrow 3D$). As shown in Tab. 3, unidirectional fusion either improves 2D metrics or 3D met-

Fusion Direction	2D Metrics		3D Metrics	
	EPE _{2D}	ACC _{1px}	EPE _{3D}	ACC _{.05}
$2D \Rightarrow 3D$	3.41	79.5%	0.036	89.6%
$2D \Leftarrow 3D$	2.25	83.9%	0.067	74.0%
$2D \Leftrightarrow 3D$	2.18	84.3%	0.033	91.4%

Table 3. Unidirectional Fusion vs. Bidirectional Fusion. Bidirectional fusion provides the best results for all metrics since both modalities can benefit each other.

rics, while bidirectional fusion provides better results for both modalities. Moreover, compared with unidirectional fusion, bidirectional fusion improves the best EPE_{2D} from 2.25 to 2.18 and EPE_{3D} from 0.036 to 0.033, suggesting that the improvement of one modality can also benefit the other.

Early/Late-Fusion vs. Multi-Stage Fusion. As mentioned above, flow estimation typically consists of several stages including feature extraction, cost volume, and feature decoding. Here, we verify the effectiveness of feature fusion for each stage, as shown in Tab. 4. The top row denotes the version where no fusion connection exists between the two branches. Both “early-fusion” and “late-fusion” (row 2, 3, 4) can only provide sub-optimal results. In contrast, fusing features at all three stages brings significant improvements compared to “early/late-fusion” (see the supplementary material for more details).

Stages			2D Metrics		3D Metrics	
P	C	D	EPE _{2D}	ACC _{1px}	EPE _{3D}	ACC _{.05}
-	-	-	3.42	79.5%	0.067	74.1%
✓	-	-	2.69	81.5%	0.047	87.6%
-	✓	-	2.42	82.3%	0.037	89.6%
-	-	✓	2.40	82.7%	0.038	88.1%
✓	✓	-	2.29	83.2%	0.034	90.7%
✓	✓	✓	2.18	84.3%	0.033	91.4%

Table 4. Early/Late-Fusion vs. Multi-Stage Fusion. P , C , D denotes features of pyramid, cost volume, and flow decoder respectively. Fusing at all three stages gives the best result.

Configurations			2D Metrics	
k -NN	2D Similarity	SOP	EPE _{2D}	ACC _{1px}
-	-	-	2.30	83.3%
$k = 1$	-	-	2.24	84.3%
$k = 1$	✓	-	2.18	84.3%
$k = 3$	✓	MEAN	2.19	84.5%
$k = 3$	✓	MAX	2.19	84.4%

Table 5. Ablations on the interpolation module of Bi-CLFM. We only report 2D metrics since 3D metrics are all similar.

Setup	IDS	2D Metrics		3D Metrics	
		EPE _{2D}	ACC _{1px}	EPE _{3D}	ACC _{.05}
C+L	-	2.24	83.2%	0.036	88.7%
C+L	✓	2.18	84.3%	0.033	91.4%
L	-	-	-	0.073	70.6%
L	✓	-	-	0.068	74.3%

Table 6. Ablations on inverse depth scaling (IDS). ‘‘L’’ denotes a LiDAR-only variation of our method where the image branch is removed, while ‘‘C+L’’ denotes our full model. IDS improves performance for both LiDAR-only and Camera-LiDAR methods.




Figure 9. Effects of introducing 2D similarities in fusion-aware interpolation, which better handles overlapping objects.

Fusion-Aware Interpolation. In Tab. 5, we test the interpolation module of Bi-CLFM with different configurations. The top row denotes a naive implementation that simply projects 3D features onto the image plane without interpolation (empty locations are filled with zeros). By introducing learnable weights into a nearest-neighbor interpolation (the 2nd row), we reduce EPE_{2D} from 2.30 to 2.24, and improve ACC_{1px} from 83.3% to 84.3%. Integrating 2D similarity measurements into the interpolation module (the 3rd row) further reduces EPE_{2D} from 2.24 to 2.18, and makes our model more robust in complex scenes with overlapping objects (as shown in Fig. 9). We also conduct two experiments by increasing k (the number of nearest neighbors) from 1 to 3, followed by a symmetric operation such

Components	Individual Timing	Total Timing
Image branch	24ms	118ms
Point branch	45ms	
Bi-CLFM (for feature pyramid)	17ms	
Bi-CLFM (for cost volume)	7ms	
Bi-CLFM (for flow decoder)	11ms	
Others	14ms	

Table 7. Breakdown of the inference time. Timing results are computed on 960x540 images with a Tesla V100 GPU.

as MEAN and MAX. However, no significant improvement is observed, suggesting that $k = 1$ is enough for interpolation.

Inverse Depth Scaling. We conduct several comparison experiments on FlyingThings3D to verify the effects of IDS. Since IDS does not require input images, we also test it on a variation of our method where the image branch is removed. As shown in Tab. 6, the performance is improved under both the camera-LiDAR and the LiDAR-only setup, suggesting that a more even distribution of points can facilitate the learning.

4.3. Analysis

Timing. In Tab. 7, we provide a breakdown of the inference time on FlyingThings3D (960x540 images). Our model takes 118ms in total with a Tesla V100 GPU.

Limitations. CamLiFlow has two limitations. First, a synchronized camera and LiDAR are required for optimal performance. If the synchronization cannot be met by some applications, our method can also take advantage of other depth sensors, such as stereo cameras (but they are less accurate and less robust than LiDAR). Another limitation is that since the two modalities are tightly coupled, the whole system will fail if one of them does not work. In the future, we plan to address this by introducing the attention mechanism to the fusion module so that the model can ‘‘choose’’ to ignore the modality if it does not work.

5. Conclusion

In this paper, we introduce CamLiFlow, a deep neural network for joint optical flow and scene flow estimation. It consists of 2D and 3D branches with multiple bidirectional connections between them in specific layers. Experiments show that CamLiFlow outperforms the previous art with fewer parameters.

Acknowledgements. This work is supported by National Natural Science Foundation of China (No.62076119, No.61921006), Program for Innovative Talents, Entrepreneur in Jiangsu Province, and Collaborative Innovation Center of Novel Software Technology and Industrialization.

References

- [1] Abhishek Badki, Orazio Gallo, Jan Kautz, and Pradeep Sen. Binary ttc: A temporal geofence for autonomous navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12946–12955, 2021. 7
- [2] Ramy Battrawy, René Schuster, Oliver Wasenmüller, Qing Rao, and Didier Stricker. Lidar-flow: Dense scene flow estimation from sparse lidar and stereo images. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7762–7769. IEEE, 2019. 2
- [3] Aseem Behl, Omid Hosseini Jafari, Siva Karthik Mustikovela, Hassan Abu Alhaja, Carsten Rother, and Andreas Geiger. Bounding boxes, segmentations and object coordinates: How important is recognition for 3d scene flow estimation in autonomous driving scenarios? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2574–2583, 2017. 1, 2, 7
- [4] Michael J Black and Paul Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer vision and image understanding*, 63(1):75–104, 1996. 2
- [5] Eric Brachmann and Carsten Rother. Neural-guided ransac: Learning where to sample model hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4322–4331, 2019. 6
- [6] Thomas Brox, Christoph Bregler, and Jitendra Malik. Large displacement optical flow. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 41–48. IEEE, 2009. 2
- [7] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision*, pages 25–36. Springer, 2004. 2
- [8] Andrés Bruhn, Joachim Weickert, and Christoph Schnörr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *International journal of computer vision*, 61(3):211–231, 2005. 2
- [9] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [10] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 2
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 6
- [12] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. Flownet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015. 2
- [13] Ziyue Feng, Longlong Jing, Peng Yin, Yingli Tian, and Bing Li. Advancing self-supervised monocular depth learning with sparse liDAR. In *5th Annual Conference on Robot Learning*, 2021. 2
- [14] Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, and Panqu Wang. Hplflownet: Hierarchical permutohedral lattice flownet for scene flow estimation on large-scale point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3254–3263, 2019. 2, 5
- [15] Yuanduo Hong, Huihui Pan, Weichao Sun, Yisong Jia, et al. Deep dual-resolution networks for real-time and accurate semantic segmentation of road scenes. *arXiv preprint arXiv:2101.06085*, 2021. 6
- [16] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [17] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 4
- [18] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8981–8989, 2018. 2
- [19] Junhwa Hur and Stefan Roth. Iterative residual refinement for joint optical flow and occlusion estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5754–5763, 2019. 2
- [20] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. 2, 5, 6
- [21] Mariano Jaimez, Mohamed Souiai, Javier Gonzalez-Jimenez, and Daniel Cremers. A primal-dual framework for real-time dense rgb-d scene flow. In *2015 IEEE international conference on robotics and automation (ICRA)*, pages 98–104. IEEE, 2015. 2
- [22] Mariano Jaimez, Mohamed Souiai, Jörg Stückler, Javier Gonzalez-Jimenez, and Daniel Cremers. Motion cooperation: Smooth piece-wise rigid scene flow from rgb-d images. In *2015 International Conference on 3D Vision*, pages 64–72. IEEE, 2015. 2
- [23] Huaizu Jiang, Deqing Sun, Varun Jampani, Zhaoyang Lv, Erik Learned-Miller, and Jan Kautz. Sense: A shared encoder network for scene-flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3195–3204, 2019. 1, 7
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [25] Yair Kittenplon, Yonina C Eldar, and Dan Raviv. Flowstep3d: Model unrolling for self-supervised scene flow estimation. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition*, pages 4114–4123, 2021. [2](#)
- [26] Congcong Li, Haoyu Ma, and Qingmin Liao. Two-stage adaptive object scene flow using hybrid cnn-crf model. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 3876–3883. IEEE, 2021. [7](#)
- [27] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 641–656, 2018. [2](#)
- [28] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. FlowNet3d: Learning scene flow in 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 529–537, 2019. [2](#), [5](#), [6](#)
- [29] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. MeteorNet: Deep learning on dynamic 3d point cloud sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9246–9255, 2019. [2](#)
- [30] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4796–4803. IEEE, 2018. [2](#)
- [31] Wei-Chiu Ma, Shenlong Wang, Rui Hu, Yuwen Xiong, and Raquel Urtasun. Deep rigid instance scene flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3614–3622, 2019. [1](#), [2](#), [7](#)
- [32] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. [2](#), [5](#), [6](#)
- [33] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. [1](#), [2](#), [5](#)
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. [5](#)
- [35] Gilles Puy, Alexandre Boulch, and Renaud Marlet. Flot: Scene flow on point clouds guided by optimal transport. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, pages 527–544. Springer, 2020. [2](#), [6](#)
- [36] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018. [2](#)
- [37] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. [2](#)
- [38] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. [2](#)
- [39] Julian Quiroga, Thomas Brox, Frédéric Devernay, and James Crowley. Dense semi-rigid scene flow estimation from rgb-d images. In *European Conference on Computer Vision*, pages 567–582. Springer, 2014. [2](#)
- [40] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4161–4170, 2017. [2](#)
- [41] Zhile Ren, Deqing Sun, Jan Kautz, and Erik Sudderth. Cascaded scene flow prediction using semantic segmentation. In *2017 International Conference on 3D Vision (3DV)*, pages 225–233. IEEE, 2017. [1](#)
- [42] Rishav Rishav, Ramy Battrawy, René Schuster, Oliver Wasenmüller, and Didier Stricker. Deeplidarflow: A deep learning architecture for scene flow estimation using monocular camera and sparse lidar. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10460–10467. IEEE, 2020. [1](#), [2](#), [6](#)
- [43] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. [2](#), [3](#), [4](#), [6](#)
- [44] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020. [2](#), [6](#)
- [45] Zachary Teed and Jia Deng. Raft-3d: Scene flow using rigid-motion embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8375–8384, 2021. [1](#), [2](#), [6](#), [7](#)
- [46] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a piecewise rigid scene model. *International Journal of Computer Vision*, 115(1):1–28, 2015. [1](#)
- [47] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. PointPainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4604–4612, 2020. [2](#)
- [48] Zirui Wang, Shuda Li, Henry Howard-Jenkins, Victor Prisacariu, and Min Chen. FlowNet3d++: Geometric losses for deep scene flow estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 91–98, 2020. [2](#)
- [49] Yi Wei, Ziyi Wang, Yongming Rao, Jiwen Lu, and Jie Zhou. Pv-raft: Point-voxel correlation fields for scene flow estimation of point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6954–6963, 2021. [2](#)
- [50] Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui, and Cordelia Schmid. DeepFlow: Large displacement optical flow with deep matching. In *Proceedings of the IEEE international conference on computer vision*, pages 1385–1392, 2013. [2](#)

- [51] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2019. 2, 4
- [52] Wenxuan Wu, Zhiyuan Wang, Zhuwen Li, Wei Liu, and Li Fuxin. Pointpwc-net: A coarse-to-fine network for supervised and self-supervised scene flow estimation on 3d point clouds. *arXiv preprint arXiv:1911.12408*, 2019. 2, 3, 4, 5, 6
- [53] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 244–253, 2018. 2
- [54] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. *Advances in neural information processing systems*, 32:794–805, 2019. 2
- [55] Gengshan Yang and Deva Ramanan. Upgrading optical flow to 3d scene flow through optical expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1334–1343, 2020. 1, 2, 6, 7
- [56] Gengshan Yang and Deva Ramanan. Learning to segment rigid motions from two frames. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1266–1275, 2021. 1, 2, 6, 7
- [57] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019. 2
- [58] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. 6