What Matters for 3D Scene Flow Network

Guangming Wang¹, Yunzhe Hu¹, Zhe Liu², Yiyang Zhou³, Masayoshi Tomizuka³, Wei Zhan³, and Hesheng Wang^{1*}

¹ Department of Automation, Key Laboratory of System Control and Information Processing of Ministry of Education, Key Laboratory of Marine Intelligent Equipment and System of Ministry of Education, Shanghai Jiao Tong University
² MoE Key Lab of Artificial Intelligence, Shanghai Jiao Tong University

³ Mechanical Systems Control Laboratory, University of California, Berkeley {wangguangming, huyz7830, liuzhesjtu, wanghesheng}@sjtu.edu.cn {yiyang.zhou, tomizuka, wzhan}@berkeley.edu

Abstract. 3D scene flow estimation from point clouds is a low-level 3D motion perception task in computer vision. Flow embedding is a commonly used technique in scene flow estimation, and it encodes the point motion between two consecutive frames. Thus, it is critical for the flow embeddings to capture the correct overall direction of the motion. However, previous works only search locally to determine a soft correspondence, ignoring the distant points that turn out to be the actual matching ones. In addition, the estimated correspondence is usually from the forward direction of the adjacent point clouds, and may not be consistent with the estimated correspondence acquired from the backward direction. To tackle these problems, we propose a novel all-to-all flow embedding layer with backward reliability validation during the initial scene flow estimation. Besides, we investigate and compare several design choices in key components of the 3D scene flow network, including the point similarity calculation, input elements of predictor, and predictor & refinement level design. After carefully choosing the most effective designs, we are able to present a model that achieves the stateof-the-art performance on FlyingThings3D and KITTI Scene Flow datasets. Our proposed model surpasses all existing methods by at least 38.2% on FlyingThings3D dataset and 24.7% on KITTI Scene Flow dataset for EPE3D metric. We release our codes at https://github.com/IRMVLab/3DFlow.

Keywords: 3D scene flow estimation, 3D PWC structure, all-to-all point mixture, point clouds, 3D deep learning.

1 Introduction

As a fundamental task in computer vision, scene flow estimation aims to estimate a 3D motion field consisting of point-wise or pixel-wise 3D displacement vectors between two consecutive frames of point clouds or images. It provides a low-level representation and understanding of the motion of dynamic objects in the scene. Many applications directly benefit from the techniques used in scene flow estimation, such as semantic segmentation [20], multi-object tracking [48,41], point cloud registration

^{*} Corresponding Author. The first two authors contributed equally.

[19,45,43], etc. The performance of scene flow estimation algorithms on point clouds has been greatly improved since deep learning was first applied in [19]. Recent studies [19,4,50,28,16,17,44,39] focus more on estimating scene flow in an end-to-end fashion from two consecutive frames of raw 3D point clouds. These approaches predict scene flow with only 3D coordinates of point clouds as inputs with no need for any prior knowledge of the scene structure. This paper also focuses on such a research topic.

Previous learning-based methods [19,50,44] adopt flow embedding to correlate adjacent frames of point clouds and to encode point motion. Their models then propagate the flow embedding through set upconv layers [19] or coarse-to-fine warping [50,44] to regress the scene flow. FlowNet3D [19], for example, introduces the flow embedding in a point-to-patch manner, which means that a specific point in the first point cloud PC_1 merely uses several neighbouring points in the second point cloud PC_2 to learn the correlation. PointPWC-Net [50] further improves it and proposes to learn a patch-topatch flow embedding, which adds a second point-to-patch embedding process in PC_1 itself after the first point-to-patch embedding between two point clouds. In addition, Pyramid, Warping, and Cost volume (PWC) structure [32] is introduced to refine the scene flow for several times. HALFlow [44] also follows this PWC structure [32] but improves by introducing the attention mechanism in both embedding processes. However, during the first embedding process between two point clouds, these methods only search for K Nearest Neighbours (KNN) in PC_2 to aggregate correspondence information. Practically, K is substantially smaller than the total number of points in the second frame, making it possible for a point in PC_1 to miss the correct yet distant matching point in PC_2 . Moreover, it is extremely important to obtain a reliable correlation when calculating it for the first time because it encodes the overall direction of the flow. The scene flow will be eventually misguided if this issue is ignored. To tackle this problem, we introduce a novel all-to-all flow embedding layer based on the double attentive flow embedding layer in HALFlow [44]. With all-to-all embedding, each point in PC_1 will use all points in PC_2 for correlation during the first embedding process, and each point in PC_2 can therefore obtain the correlation with all points in PC_1 too. This mechanism allows that the feature correlation of all points can be exhaustively utilized from both sides and reliable correspondence estimation can be further achieved.

This all-to-all mechanism, however, cannot guarantee that the reliable correlation is bi-directional. That is, the estimated match pair for a specific point in PC_1 in the forward direction may not be consistent with the match for the corresponding matched point in PC_2 in the backward direction. Therefore, we need another constraint on the backward match to validate its consistency with the forward pass. Mittal *et al.* [26] utilize a similar mechanism by designing a cycle-consistency loss to achieve selfsupervising the scene flow estimation. However, we expect to directly incorporate this constraint of backward validation into our network to allow the network to learn this ability in forward reasoning. To this end, we propose backward reliability validation, a joint learning method of backward constraint in the all-to-all flow embedding layer.

Furthermore, there are several components of our network with alternative designs either from themselves or from other works that could affect the performance. Therefore, we conduct a series of ablation studies to compare different designs and to explore which elements are important and which designs are suitable for 3D scene flow network, including the point similarity calculation, predictor elements choice, and predictor & refinement level design. Our key contributions are as follows:

- A novel all-to-all point mixture module with backward reliability validation is proposed for reliable correlation between point clouds. The all-to-all mechanism is adopted to capture reliable match candidates from the distance, and backward information is integrated in the inference process to validate the matching consistency.
- Different designs and techniques of 3D scene flow network are widely compared and analyzed. *Point Similarity Calculation, Designs of Scene Flow Predictor, Input elements of Scene Flow Predictor*, and *Flow Refinement Level Design* are individually discussed and evaluated to showcase what matters in 3D scene flow network.
- 3. Experiments demonstrate that our model achieves state-of-the-art performance, reducing EPE3D metric by at least 38.2% on FlyingThings3D dataset [23] and 24.7% on KITTI Scene Flow dataset [25]. The effectiveness of proposed techniques and choices of network designs are demonstrated through extensive ablation studies.

2 Related Work

The concept of scene flow is first introduced by Vedula *et al.* [36] as the 3D motion field in real-world scenarios. Many previous works estimate scene flow by recovering the 3D motion from optical flow and depth information on 2D image pairs, either using RGBstereo [7,27,35,24,37,38,23,22] or RGB-D [5,6,11] data. There has also been some recent works focusing on recovering scene flow from monocular camera [52,8,53,9,42]. However, since scene flow indicates the 3D motion, directly estimating scene flow from 3D data input can enable direct optimization and higher accuracy. The applications of LiDARs in recent years have created more available raw data of point cloud, and pointcloud-based scene flow estimation approaches [2,34] are rapidly emerging.

Since deep learning has shown excellent performance for raw point-cloud-based tasks [40,46,21,51] with the proposal of PointNet [29] and PointNet++ [30], many works estimate scene flow directly from raw point clouds in an end-to-end fashion. FlowNet3D [19] presents the first end-to-end scene flow estimation framework on point clouds. It uses PointNet++ [30] to extract local point features and introduces a flow embedding layer to encode the point motions. HPLFlow-Net [4] leverages the idea from Bilateral Convolutional Layers (BCL) [14,12] and proposes DownBCL, UpBCL, and CorrBCL designs to restore structural information of large-scale point clouds.

More recent works focus on improving the network performance through introducing new techniques or incorporating new components. FLOT [28] proposes to find the correspondences from an optimal transport module by graph matching. PointPWC-Net [50] follows a coarse-to-fine fashion for scene flow estimation on point clouds. It extends the important component of cost volume in optical flow network [32] and proposes a novel point-based patch-to-patch cost volume. HALFlow [44] improves the aforementioned cost volume by a novel double attentive flow embedding method that distributes more weights on task-related regions. HCRF-Flow [17] focuses on maintaining the local geometric smoothness with the help of Conditional Random Fields (CRFs) in deep neural networks and proposes a high-order CRFs module as the formulations of spatial smoothness and rigid motion constraints. FESTA [47] improves naive Farthest



Fig. 1. The detailed architecture of our network. Three set conv layers for PC_1 and four set conv layers for PC_2 constitute the hierarchical point feature abstraction module. The all-to-all point mixture module consists of one layer of all-to-all flow embedding followed by two set conv layers. Four flow refinement layers are constructed in the hierarchical flow refinement module.

Point Sampling (FPS) by proposing a trainable Aggregate Pooling (AP) to adaptively shift points to invariant positions. Inspired by [28], FlowStep3D [16] designs a Global Correlation Unit that computes a soft correlation matrix to guide the initially estimated scene flow and adopts Gated Recurrent Unit (GRU) for local flow update. PV-RAFT [49] leverages point-based and voxel-based features and presents point-voxel correlation fields to capture both local and long-range dependencies for point pairs.

In [50,16], different information is utilized for updating local scene flow but they do not show which information is more important for the input of the updating unit. In [16,49], GRU is used for iterative flow update inspired by RAFT [33] and claimed to be more effective than a fully-connected layer. Since RAFT [33] has shown a promising performance on optical flow, we also want to know whether the use of GRU will improve the performance in other scene flow network structures. In [28,16], element-wise product and cosine similarity are used to represent correlation between points while concatenation of point feature is implemented for learning correlation in [44], but none of them gives evaluation about which one is better. This paper will discuss the above issues and compare what matters for 3D scene flow network based on PWC structure.

3 3D Scene Flow Network

3.1 Network Architecture

Our proposed network, illustrated in Fig. 1, takes in two frames of point clouds with 4N points in each, which are PC_1 and PC_2 , and estimates N points' scene flow from coarse to fine. Our network is comprised of three modules: 1) Hierarchical Point Feature Abstraction, 2) All-to-All Point Mixture, and 3) Hierarchical Flow Refinement.

The hierarchical point feature abstraction module has three set conv layers from [19] for PC_1 and four set conv layers for PC_2 . Each set conv layer performs down-sampling

operation on the input points and extracts local features of the down-sampled points. The same level of the set conv layers shares the same weights. Then, the proposed all-toall flow embedding layer correlates two point clouds and learns the flow embedding. We then use two set conv layers after the flow embedding layer for smoothness. Next, the output of the all-to-all point mixture module is up-sampled by the set upconv layer from [19] to generate the initial flow embedding. A Fully-Connected (FC) layer is thereafter applied on the initial flow embedding are both fed into the hierarchical flow refinement module and refined iteratively to derive the final scene flow using the information from specific level. The skip connections indicate which level of information is utilized.

3.2 Hierarchical Point Feature Abstraction

In the hierarchical point feature abstraction module, two consecutive point clouds are down-sampled and encoded through a series of set conv layers respectively. We adopt the set conv layer in PointNet++ [30] to perform point feature abstraction.

Each set conv layer consumes n points $\{(x_i, p_i) \mid i = 1, ..., n\}$, where $x_i \in \mathbb{R}^3$ and $p_i \in \mathbb{R}^c$ represent 3D coordinate and the point feature. The output of each layer are n' sampled point $\{(x'_j, p'_j) \mid j = 1, ..., n'\}$ with $x'_j \in \mathbb{R}^3$ and $p'_j \in \mathbb{R}^{c'}$ denoting the 3D coordinate and extracted local feature. All of the output n' (n' < n) are sampled from the input n points using Farthest Point Sampling (FPS) [30].

For each of the n' sampled point, its K nearest neighbours $\{(x_i^k, p_i^k) | k = 1, ..., K\}$ are selected from the input n points. Then, a learnable shared Multi-Layer Perceptron (MLP) and max-pooling operation are adopted to extract the point feature p'_j of each sampled point from the K neighbouring points. The point feature p'_j is formulated as:

$$p'_{j} = \underset{k=1,\ldots,K}{\operatorname{maxpool}}(\operatorname{MLP}((x^{k}_{j} - x'_{j}) \oplus p^{k}_{j})), \tag{1}$$

where \oplus indicates concatenation operation. *maxpool* means max-pooling operation.

3.3 All-to-All Point Mixture

The inputs of the all-to-all embedding layer are two consecutive frames of point clouds: $PC_1 = \{(x_i, p_i) \mid i = 1, ..., n_1\}$ and $PC_2 = \{(y_j, q_j) \mid j = 1, ..., n_2\}$, sampled in the hierarchical point feature abstraction module. $x_i, y_j \in \mathbb{R}^3$ indicate 3D coordinates and $p_i, q_j \in \mathbb{R}^c$ indicate the point feature. The output of the layer will be the flow embedding $E = \{e_i \mid e_i \in \mathbb{R}^c, i = 1..., n_1\}$, which utilizes exhaustive information in two point clouds and encodes motion for points in PC_1 .

Our all-to-all embedding includes a two-stage attention-based embedding process with an improved first embedding stage. For the first embedding process, we correlate points in two point clouds by incorporating the all-to-all mechanism with backward reliability validation as shown in Fig. 2. Instead of choosing only $K(K < n_2)$ nearest points, each point in PC_1 selects all n_2 points $Q_i = \{(y_i^k, q_i^k) \mid k = 1, ..., n_2\}$ from PC_2 . In this process, all n_2 points Q_i from PC_2 are utilized to embed point motion into points in PC_1 . The motion embeddings will then be updated by carefully designed



Fig. 2. The detailed calculation of the first flow embedding between PC_1 and PC_2 .

attentive weighting to derive the first flow embedding $FE = \{fe_i \mid i = 1, ..., n_1\}$. The calculation details are elaborated as below.

A 10-dimensional vector capturing the 3D Euclidean space information is first calculated as follows:

$$d_i^k = x_i \oplus y_i^k \oplus (x_i - y_i^k) \oplus \left\| x_i - y_i^k \right\|_2, \tag{2}$$

where $\|\cdot\|_2$ indicates the L_2 norm. Then, to realize backward reliability validation, we first formulate a vector that represents a form of similarity between two point clouds by applying element-wise product of PC_1 point feature p_i and PC_2 point feature q_i^k . Maxpooling operation is then performed over n_1 candidate backward embedding features, selecting the most potential and reliable matching candidates in PC_1 for each point of PC_2 . An FC layer is then adopted to encode the backward reliability information. The calculation of backward validation vector is as follows:

$$s_i^k = \mathsf{FC}(\underset{i=1,\dots,n_1}{\operatorname{maxpool}}(p_i \odot q_i^k)), \tag{3}$$

where \odot denotes dot product. The first flow embedding before attentive weighting is then formulated as:

$$h_i^k = \mathrm{MLP}(d_i^k \oplus p_i \oplus q_i^k \oplus s_i^k).$$
(4)

Specifically, p_i and q_i^k are normalized on the feature channel before concatenation. Given the 10-dimensional vector d_i^k , the first attentive weights for soft aggregation of the queried points can be written as:

$$w_i^k = \underset{k=1,\dots,n_2}{\operatorname{softmax}} (\operatorname{MLP}(\operatorname{FC}(d_i^k) \oplus h_i^k)).$$
(5)

The first flow embedding $FE = \{fe_i \mid i = 1, ..., n_1\}$ corresponding to points with x_i coordinates is calculated as:

$$fe_i = \sum_{k=1}^{n_2} h_i^k \odot w_i^k.$$
(6)

For the second flow embedding process, we follow the same process as [44], which is an aggregation process within the PC_1 self with attention. Each point in PC_1 selects

several nearest neighbours in PC_1 self, and the neighbourhood flow embeddings in FE will be aggregated into each point in PC_1 to obtain the second flow embedding $E = \{e_i \mid i = 1, ..., n_1\}$, which is the output of the all-to-all flow embedding layer.

Calculation of Point Similarity In formula (4), point feature p_i of PC_1 and q_i^k of PC_2 are concatenated to learn the similarity between points from two point clouds. However, there are other ways [28,16,13] to calculate and represent the point similarity: 1) product similarity: the direct dot product of p_i and q_i^k as $sim(p_i, q_i^k) = \langle p_i, q_i^k \rangle$, 2) cosine product similarity: the dot product of p_i and q_i^k divided by their respective L_2 norm as $sim(p_i, q_i^k) = \langle \frac{p_i}{\|p_i\|_2}, \frac{q_i^k}{\|q_i^k\|_2} \rangle$, and 3) normalized product similarity: the dot product of p_i and q_i^k normalized by their respective mean value μ and standard deviation σ over each feature dimension as $sim(p_i, q_i^k) = \langle \frac{p_i - \mu_i}{\sigma_i}, \frac{q_i^k - \mu_i^k}{\sigma_i^k} \rangle$. We intend to explore whether the concatenation of point feature is more suitable to represent similarity in our network compared with the product similarity presented above. By replacing the concatenation of feature with different forms of product of feature in formula (4), the effectiveness of our design is demonstrated in experiments.

3.4 Hierarchical Flow Refinement

The hierarchical flow refinement module consists of four flow refinement layers. The layer takes coarse sparse flow and coarse sparse flow embedding as inputs with information of PC_1 and PC_2 from the previous level while producing refined flow and refined flow embedding as outputs, as illustrated in Fig. 3. It contains four main components: 1) Set Upconv Layer, 2) Position Warping Layer, 3) Attentive Flow Re-embedding Layer, and 4) Scene Flow Predictor. For the first flow refinement layer, the set upconv layer is eliminated to keep the point number unchanged for suitable multi-scale supervision.

Set Upconv Layer In order to up-sample the coarse sparse flow embedding, the set upconv layer in [19] is adopted here to propagate flow embedding from sparse level to dense level. The inputs of this layer are n points with coarse sparse flow embedding $\{(x_i, se_i) \mid se_i \in \mathbb{R}^{d_{sparse}}, i = 1, ..., n\}$ and n' (n' > n) points with feature $\{(x'_j, p'_j) \mid j = 1, ..., n'\}$ from the previous dense level. The outputs are n' points with dense flow embedding $\{(x'_j, de_j) \mid de_j \in \mathbb{R}^{d_{dense}}, j = 1, ..., n'\}$. Specifically, each of the n' dense points will select its KNN from the sparse n points, and the coarse sparse flow embedding by MLP.

Position Warping Layer As a coarse-to-fine style, coarse dense flow $\{f_i^{dense} \mid i = 1, \ldots, n_1\}$ is first obtained from coarse sparse flow through Three-Nearest Neighbours (Three-NN) interpolation. Next, the coordinates of the first point cloud $PC_1 = \{(x_i, p_i) \mid i = , ..., n_1\}$ are updated by warping PC_1 with coarse dense flow. The warped PC_1 is signified as $PC'_1 = \{(x'_i, p_i) \mid i = 1, ..., n_1\}$, where $x'_i = x_i + f_i^{dense}$.

Attentive Flow Re-embedding Layer The attentive flow embedding layer proposed in [44] is applied to derive a new flow re-embedding $\{re_i \mid re_i \in \mathbb{R}^{d_{re}}, i = 1, \ldots, = n_1\}$ between PC'_1 and PC_2 . Here, the flow re-embedding contains the flow encoding from each point in PC'_1 toward PC_2 , which is essential in the subsequent refinement.

Scene Flow Predictor The scene flow predictor aims to refine coarse dense flow embedding for the input of later flow refinement layer. It takes five elements as inputs: 1) the up-sampled coarse dense flow embedding $de_i \in \mathbb{R}^{d_{dense}}$, 2) the



Fig. 3. The details of flow refinement layer.

flow re-embedding $re_i \in \mathbb{R}^{d_{re}}$, 3) the point feature of the first point cloud $p_i \in \mathbb{R}^{d_{pc_1}}$, 4) the coarse dense flow $f_i^{dense} \in \mathbb{R}^3$, and 5) the dense flow feature $f_i^{enc.} \in \mathbb{R}^{d_{enc.}}$. Specifically, the coarse dense flow is encoded by two set conv layers [19] to derive the dense flow feature $f_i^{enc.}$, but the number of points remain unchanged instead of being down-sampled. The refined dense flow embedding is formulated as:

$$de'_{i} = \mathsf{MLP}(de_{i} \oplus re_{i} \oplus p_{i} \oplus f_{i} \oplus f_{i}^{enc.}).$$

$$\tag{7}$$

Finally, we adopt a residual flow learning structure to estimate the refined scene flow. To be specific, an FC layer is first applied on the refined flow embedding to produce the residual flow f_i^{res} . Next, the refined scene flow is generated by adding f_i^{res} to f_i^{dense} . The calculation of refined scene flow is as follows:

$$f_i^{res} = FC(de_i'), \tag{8}$$

$$f_i = f_i^{dense} + f_i^{res}.$$
(9)

Designs of Scene Flow Predictor: The scene flow predictor corrects the coarse scene flow by regressing residual flow from the refined flow embedding. In this paper, we adopt the concatenation of all five inputs of scene flow predictor and directly feed it into shared MLP to derive the refined flow embedding. On the other hand, FlowStep3D [16] and PV-RAFT [49] propose to use a GRU-based gated activation unit on point clouds, inspired by RAFT [33], for updating a hidden state. Given a hidden state $h_{l-1} \in \mathbb{R}^c$ from previous iteration and a current iteration vector x_l , h_{l-1} is updated as follows:

$$z_l = \sigma(\operatorname{SetConv}_z(h_{l-1} \oplus x_l)), \tag{10}$$

$$r_l = \sigma(\operatorname{SetConv}_{\mathbf{r}}(h_{l-1} \oplus x_l)), \tag{11}$$

$$h_{l} = \tanh(\operatorname{SetConv}_{h}((r_{l} \odot h_{l-1}) \oplus x_{l})), \qquad (12)$$

$$h_{l} = (1 - z_{l}) \odot h_{l-1} + z_{l} \odot h_{l},$$
(13)

where $\sigma(\cdot)$ represents sigmoid activation function.

In particular, $x_l \in \mathbb{R}^{d_{re}+d_{pc_1}+d_{enc.}+3}$ is defined as the concatenation of flow reembedding, the point feature of PC_1 , the coarse dense flow, and the dense flow feature. We refer to the dense flow embedding as the hidden state that will be refined iteratively. Then, we can consider replacing our scene flow predictor with this newly designed GRU-based updating unit. Since FlowStep3D [16] and PV-RAFT [49] claim that this GRU-based updating mechanism outperforms the fully-connected structure which is implemented in MLP, we will validate the performance of our scene flow predictor in our network architecture compared with this GRU-based method in the experiment.

Input of Scene Flow Predictor: Another issue we want to investigate is what information is needed to predict the finer flow embedding. i.e. what information is needed in the input of the scene flow predictor. PointPWC-Net [50] uses the flow re-embedding, the point feature of PC_1 , the coarse dense flow, and the up-sampled coarse dense flow embedding as inputs. FlowStep3D [16] additionally includes the dense flow feature but do not add the up-sampled flow embedding. PV-RAFT [49], HCRF-Flow [17], and FESTA [47] all only use the flow re-embedding, the point feature of PC_1 , and the coarse flow. HALFlow [44] does not add the dense flow feature. In this sense, we explore whether removing certain elements from the input will degrade the performance and the extent to which those information contributes to the performance. We consider the flow re-embedding as an indispensable element because it encodes the motion between warped PC_1 and PC_2 in the current level.

Level of Flow Refinement Layer: In [44], three flow refinement layers are applied for a more-for-less network architecture, which estimates N points' scene flow from 4N points of raw input. Only Three-NN interpolation is used for up-sampling in the finest level of flow estimation. Apparently, this structure does not exhaustively leverage the functionality of the whole components in the flow refinement layer. Therefore, we choose to employ an additional flow refinement layer on the finest level to estimate the final scene flow. The above consideration raises another question: Can the performance be further improved if more information from denser level is taken into account? Here we consider using raw point cloud with 4N points to estimate the 4N points' scene flow. The best choice will be demonstrated in the experiment.

4 Experiments

4.1 Datasets and Data Preprocess

Because of the inherent difficulty in acquiring large-scale ground-truth scene flow of the real world, we resort to the common synthetic FlyingThings3D dataset [23] for training and evaluation. We first train our network on FlyingThings3D dataset [23], and then directly test our trained model on real-world LiDAR scans from KITTI scene flow dataset [25] without any fine-tuning to demonstrate the generalization capability.

There are two common versions of preparing point clouds from FlyingThings3D dataset [23] and KITTI scene flow dataset [25]. The first version of data preprocessing is proposed by HPLFlowNet [4] and adopted in [4,50,44,17,16]. It does not contain occlusion for input point clouds, which means each point in PC_1 has its corresponding point in PC_2 when warped by its ground-truth scene flow. The second version is

proposed by FlowNet3D [19]. This version provides occluded point clouds as inputs and masks that indicate the invalid points without corresponding ones in the adjacent frame. These masks are also used in computing training loss and evaluation metrics. To compare with all 3D scene flow estimation methods to our knowledge, we follow two versions of data preprocessing and conduct experiments on both versions of datasets. More details about datasets and preprocessing is in the supplementary material.

4.2 Training and Evaluation Details

Training Loss We train our network in a supervised manner at different levels, similar to [32,50]. Suppose the predicted scene flow of each point at level l is $\{f_i^l \in \mathbb{R}^3 \mid i = 1, ..., N_l\}$ and the ground-truth scene flow is $\{GT(f_i^l) \in \mathbb{R}^3 \mid i = 1, ..., N_l\}$. Here, N_l denotes the number of points at level l. Our training loss can be therefore written as:

$$Loss = \sum_{l=1}^{4} \psi_l \frac{1}{N_l} \sum_{i=1}^{N_l} \left\| f_i^l - GT(f_i^l) \right\|_2, \tag{14}$$

where ψ_l indicates the weight at level l. We define the finest level, which is also the level with the densest points, as level l = 1. Specifically, our network takes in 4N = 8,912 as inputs, $N_1 = N = 2,048$, $N_2 = N/2 = 1,024$, $N_3 = N/8 = 256$, and $N_4 = N/32 = 64$. The loss weights are $\psi_1 = 0.02$, $\psi_2 = 0.04$, $\psi_3 = 0.08$, and $\psi_4 = 0.16$.

Implementation Details For the training and evaluation process of our network, 8,192 points are randomly sampled as inputs from the raw points clouds of two consecutive frames. Only 3D XYZ coordinates of the point clouds are fed into our network, like [4,50,44,17,16]. For fair comparison with previous methods, on FlyingThings3D prepared by [4], we first train our network on one quarter of the training set (4,910 pairs) and then fine-tune our model on the complete training set to speed up the training process. On FlyingThings3D prepared by [19], we train our model without fine-tuning.

We conduct all the experiments on a single Titan RTX GPU. Pre-training is done for 800 epochs, and fine-tuning lasts for 200 epochs after loading the pre-trained weights. Batch size is 14. Adam optimizer [15] is used in training, and $\beta_1 = 0.9$, $\beta_2 = 0.99$. The initial learning rate is 0.001 and decays for every 80 epochs exponentially with decay rate $\gamma = 0.5$. Our supplementary material provides all details of network parameters.

Evaluation Metrics We adopt the same evaluation metrics used in [4,50,44,17] to evaluate our model for fair comparison, including EPE3D(m), Acc3D Strict, Acc3D Relax, Outliers3D, EPE2D(px), and Acc2D. The detailed description of the metrics is shown in the supplementary material.

5 Results

5.1 Comparison with State-of-the-Art (SOTA)

Table 1 shows the quantitative comparison between previous state-of-the-arts and our approach on FlyingThings3D dataset [23] and KITTI scene flow dataset [25] prepared

Table 1. The quantitative comparison between recent state-of-the-art methods and ours on FlyingThings3D and KITTI scene flow datasets prepared by Gu *et al.* [4] without occlusion. All listed approaches are only trained on FlyingThings3D dataset. The best results are in bold. "Full" means fully-supervised training.

Evaluation Dataset	Method	Training Data	Input	Sup.	EPE3D	Acc3D Strict	Acc3D Relax	Outliers3D	EPE2D	Acc2D
	FlowNet3 [10]	Quarter	RGB stereo	Full	0.4570	0.4179	0.6168	0.6050	5.1348	0.8125
FlyingThings 3D	ICP [1]	No	Points	Full	0.4062	0.1614	0.3038	0.8796	23.2280	0.2913
	FlowNet3D [19]	Quarter	Points	Full	0.1136	0.4125	0.7706	0.6016	5.9740	0.5692
	SPLATFlowNet [31]	Quarter	Points	Full	0.1205	0.4197	0.7180	0.6187	6.9759	0.5512
	HPLFlowNet [4]	Quarter	Points	Full	0.0804	0.6144	0.8555	0.4287	4.6723	0.6764
	HPLFlowNet [4]	Complete	Points	Full	0.0696	—	—	—	—	_
	PointPWC-Net [50]	Complete	Points	Full	0.0588	0.7379	0.9276	0.3424	3.2390	0.7994
uataset [23]	HALFlow [44]	Quarter	Points	Full	0.0511	0.7808	0.9437	0.3093	2.8739	0.8056
	HALFlow [44]	Complete	Points	Full	0.0492	0.7850	0.9468	0.3083	2.7555	0.8111
	FLOT [28]	Complete	Points	Full	0.0520	0.7320	0.9270	0.3570	_	_
	HCRF-Flow [17]	Quarter	Points	Full	0.0488	0.8337	0.9507	0.2614	2.5652	0.8704
	PV-RAFT [49]	Complete	Points	Full	0.0461	0.8169	0.9574	0.2924	_	_
	FlowStep3D [16]	Complete	Points	Full	0.0455	0.8162	0.9614	0.2165	_	_
	Ours	Quarter	Points	Full	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108
	Ours	Complete	Points	Full	0.0281	0.9290	0.9817	0.1458	1.5229	0.9279
	FlowNet3 [10]	Quarter	RGB stereo	Full	0.9111	0.2039	0.3587	0.7463	5.1023	0.7803
	ICP [1]	No	Points	Full	0.5181	0.0669	0.1667	0.8712	27.6752	0.1056
	FlowNet3D [19]	Quarter	Points	Full	0.1767	0.3738	0.6677	0.5271	7.2141	0.5093
	SPLATFlowNet [31]	Quarter	Points	Full	0.1988	0.2174	0.5391	0.6575	8.2306	0.4189
KITTI dataset [25]	HPLFlowNet [4]	Quarter	Points	Full	0.1169	0.4783	0.7776	0.4103	4.8055	0.5938
	HPLFlowNet [4]	Complete	Points	Full	0.1113	_	_	_	_	_
	PointPWC-Net [50]	Complete	Points	Full	0.0694	0.7281	0.8884	0.2648	3.0062	0.7673
	HALFlow [44]	Quarter	Points	Full	0.0692	0.7532	0.8943	0.2529	2.8660	0.7811
	HALFlow [44]	Complete	Points	Full	0.0622	0.7649	0.9026	0.2492	2.5140	0.8128
	FLOT [28]	Complete	Points	Full	0.0560	0.7550	0.9080	0.2420	_	_
	HCRF-Flow [17]	Quarter	Points	Full	0.0531	0.8631	0.9444	0.1797	2.0700	0.8656
	PV-RAFT [49]	Complete	Points	Full	0.0560	0.8226	0.9372	0.2163	_	_
	FlowStep3D [16]	Complete	Points	Full	0.0546	0.8051	0.9254	0.1492	_	_
	Ours	Quarter	Points	Full	0.0332	0.8931	0.9528	0.1690	1.2186	0.9373
	Ours	Complete	Points	Full	0.0309	0.9047	0.9580	0.1612	1.1285	0.9451

by [4]. It is demonstrated that our approach outperforms all other methods by a large margin for both 3D and 2D metrics on FlyingThings3D dataset [23]. Meanwhile, our method also achieves the best generalization results on KITTI scene flow [25]. Specifically, we surpasses the SOTA method, FlowStep3D [16], by 38.2% for EPE3D metric on FlyingThings3D dataset [23], and 43.4% on KITTI scene flow [25] dataset.

The recent work, FESTA [47], is only tested on the datasets prepared by [19]. To compare with all the methods to our knowledge, we also present the evaluation results on FlyingThings3D dataset [23] and KITTI scene flow dataset [25] prepared by [19] in Table 2. It can be demonstrated that our approach still outperforms previous methods substantially for all 3D metrics on both datasets. Specifically, we surpasses the SOTA method, FESTA [47], by 43.2% with respect to EPE3D metric on FlyingThings3D dataset [23], and 24.7% on KITTI scene flow [25] dataset. We believe the superior performance of our method on the datasets with occlusion partly from our backward validation, which can be aware of the occlusion in the network inference.

We also present detailed visualization of the accuracy of the estimated scene flow by our approach in Fig. 4, compared with methods in [19,44]. It can be seen that our method can better handle the structures with repetitive patterns and large motions.

Table 2. Evaluation results on FlyingThings3D and KITTI Scene Flow datasets prepared by Liu *et al.* [19] with occlusion. The best results are in bold. "Self" means self-supervised training. "Full" means fully-supervised training. All fully-supervised approaches are trained on FlyingThings3D dataset. Self-Point-Flow [18] is trained on raw LiDAR data from KITTI dataset [3].

Evaluation Dataset	Method	Input	Sup.	EPE3D	Acc3D Strict	Acc3D Relax	Outliers
	FlowNet3D [19]	Points	Full	0.169	0.254	0.579	0.789
FlyingThings 3D	FLOT [28]	Points	Full	0.156	0.343	0.643	0.700
dataset [23]	FESTA [47]	Points	Full	0.111	0.431	0.744	_
	Ours	Points	Full	0.063	0.791	0.909	0.279
	Self-Point-Flow [18]	Points	Self	0.105	0.417	0.725	0.501
<i>V</i> ITTI	FlowNet3D [19]	Points	Full	0.173	0.276	0.609	0.649
KIIII dataaat [25]	FLOT [28]	Points	Full	0.110	0.419	0.721	0.486
uataset [23]	FESTA [47]	Points	Full	0.097	0.449	0.833	_
	Ours	Points	Full	0.073	0.819	0.890	0.261

5.2 Ablation Study

In this section, we demonstrate the effectiveness of the proposed all-to-all point mixture module and validate our network design choices compared with other discussed alternative structures through a series of ablation studies. All methods in the ablation studies are trained on $\frac{1}{4}$ of the training set (4,910 pairs) of FlyingThings3D dataset [23] prepared by [4] and evaluated using the corresponding evaluation set.

All-to-all and Backward Information In order to demonstrate the effectiveness of both the all-to-all mechanism and the backward validation vector, we first remove the vector that stores the backward information in the all-to-all flow embedding layer, which means s_i^k is removed from formula (4). Then, the all-to-all mechanism is removed, which means points from PC_1 will only select K nearest neighbours instead. The results in Table 3(a) show that both the all-to-all mechanism and the backward validation vector contribute to the improvement of performance. In fact, the all-to-all mechanism enables the querying point in PC_1 to expand its searching range from K nearest neighbours to all points in PC_2 to determine its most reliable matching candidate. The backward validation vector acquired based on the all-to-all mechanism brings in the information that imposes bi-directional consensus from the backward direction. Therefore, the network can be guided by this backward validation information to better learn the correct matching and estimate more accurate scene flow.

Calculation of Point Similarity Three different forms of similarity calculation are discussed in section 3.3. Since point similarity design can affect the correlation of adjacent point clouds to a large extent, we compare these product similarity methods with ours. The results in Table 3(b) demonstrate that the concatenation of point feature has the best performance than product forms of point similarity. Representing point similarity via concatenation is more suited in our network. We believe this is because the concatenation operation allows the network to fully exploit its self-learning ability.



Fig. 4. The visualization results of the proposed method, compared with FlowNet3D [19] and HALFlow [44], on FlyingThings3D (left) and KITTI scene flow (right) datasets prepared by Gu *et al.* [4]. Blue points indicate PC_1 . Green points indicate accurate predictions $\widetilde{PC_2} = PC_1 + F$ and red points indicate inaccurate predictions (measured by Acc3D Relax).

Designs of Scene Flow Predictor Since our scene flow predictor is only implemented with MLP, we want to know whether GRU in [16,49] can improve the performance. We consider the flow embedding as the hidden state to be updated and replace our MLP based scene flow predictor with GRU. Table 3(c) show that the performance actually degrades using GRU. For our coarse-to-fine network, GRU is less suitable for updating and refining the flow embedding. We believe this is due to the difference of the point number at different levels, which is different from FlowStep3D [16] and PV-RAFT [49].

Input of Scene Flow Predictor As the scene flow predictor serves to refine the flow embedding for regression of more accurate scene flow, we are interested in what matters in the input information of the scene flow predictor and to what extent each of the inputs contributes to the performance. Therefore, we respectively ablate each of the five inputs except the flow re-embedding in our scene flow predictor to demonstrate their importance. Table 3(d) show that the up-sampled flow embedding is the most important element of inputs, which is intuitive because the up-sampled flow embedding is refined in the previous layer and contains abundant motion information. The removal of the coarse flow and flow feature also causes a slight decline in quantitative performance, which demonstrates that they can also provide some instructions for the improvement of the refinement.

Level of Flow Refinement Layer In our hierarchical flow refinement module, four flow refinement layers exhaustively utilize all four levels of sampled point feature, in-

	Method	EPE3D	Acc3D Strict	Acc3D Relax	Outliers	EPE2D	Acc2D
(a)	Ours w/o backward validation	0.0332	0.9044	0.9743	0.1766	1.8221	0.9065
	Ours w/o backward validation and all-to-all mechanism	0.0349	0.9001	0.9725	0.1798	1.9819	0.9032
	Ours (full, with backward validation and all-to-all mechanism)	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108
(b)	Ours (with product similarity)	0.0356	0.8872	0.9692	0.1953	1.9872	0.8870
	Ours (with cosine product similarity)	0.0370	0.8755	0.9670	0.2142	2.0637	0.8746
	Ours (with normalized product similarity)	0.0339	0.8971	0.9724	0.1845	1.8790	0.8965
	Ours (full, with concatenated similarity)	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108
(c)	Ours (replace Scene Flow Predictor with GRU)	0.0350	0.8892	0.9668	0.1827	1.9274	0.8896
	Ours (full, with Scene Flow Predictor)	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108
(d)	Ours w/o features of PC_1 in Scene Flow Predictor	0.0333	0.9047	0.9743	0.1740	1.8428	0.9073
	Ours w/o up-sampled flow embedding in Scene Flow Predictor	0.0380	0.8732	0.9642	0.2099	2.0953	0.8785
	Ours w/o coarse flow in Scene Flow Predictor	0.0323	0.9076	0.9750	0.1717	1.7760	0.9083
	Ours w/o flow feature in Scene Flow Predictor	0.0327	0.9061	0.9748	0.1740	1.8063	0.9074
	Ours (full, with complete five inputs in Scene Flow Predictor)	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108
(e)	Ours (with interpolation estimating 2048 points' flow)	0.0359	0.8844	0.9691	0.2004	1.9511	0.8911
	Ours (with interpolation estimating 8192 points' flow)	0.0332	0.9043	0.9739	0.1740	1.8039	0.9076
	Ours (full, with flow refinement estimating 2048 points' flow)	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108

Table 3. The ablation study results on FlyingThings3D dataset prepared by Gu et al. [4].

dicated by skip connection. Compared with [44], we add an additional flow refinement layer on the finest level. In order to validate its effectiveness, we remove this flow refinement layer and instead use Three-NN interpolation to obtain the final N points' scene flow from 4N points' inputs, as [44]. In addition, we further apply the Three-NN interpolation to obtain 4N points' scene flow out of 4N points' inputs to see whether more information from raw input will result in an improvement. It turns out in Table 3(e) that both of these designs will degrade the network performance and our four layers of flow refinement that leverage all levels' sampled point feature prove the most suitable and effective structure design for our network.

6 Conclusion

In this paper, a novel all-to-all point mixture module with backward reliability validation is proposed for reliable correlation. In addition, different designs and techniques for key components of our network are thoroughly compared. We provide a series of ablation studies to show the contributions of each element in key components and to demonstrate what matters in scene flow network. Quantitative results on FlyingThings3D [23] and KITTI scene flow dataset [25] show that our method achieves SOTA performance. Our comparison and analysis on design choices of key components and structure are expected to facilitate the design of scene flow network in future research.

Acknowledgement. This work was supported in part by the Natural Science Foundation of China under Grant 62073222, Grant U21A20480, and Grant U1913204; in part by the Science and Technology Commission of Shanghai Municipality under Grant 21511101900; and in part by the Open Research Projects of Zhejiang Laboratory under Grant 2022NB0AB01. The authors gratefully appreciate the contribution of Chaokang Jiang from China University of Mining and Technology, and Xinrui Wu from Shanghai Jiao Tong University.

References

- Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: Sensor fusion IV: control paradigms and data structures. vol. 1611, pp. 586–606. Spie (1992)
- Dewan, A., Caselitz, T., Tipaldi, G.D., Burgard, W.: Rigid scene flow for 3d lidar scans. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1765–1770. IEEE (2016)
- Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. The International Journal of Robotics Research 32(11), 1231–1237 (2013)
- Gu, X., Wang, Y., Wu, C., Lee, Y.J., Wang, P.: 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. pp. 3254–3263 (2019)
- Hadfield, S., Bowden, R.: Kinecting the dots: Particle based scene flow from depth sensors. In: 2011 International Conference on Computer Vision. pp. 2290–2295. IEEE (2011)
- Herbst, E., Ren, X., Fox, D.: Rgb-d flow: Dense 3-d motion estimation using color and depth. In: 2013 IEEE international conference on robotics and automation. pp. 2276–2282. IEEE (2013)
- Huguet, F., Devernay, F.: A variational method for scene flow estimation from stereo sequences. In: 2007 IEEE 11th International Conference on Computer Vision. pp. 1–7. IEEE (2007)
- Hur, J., Roth, S.: Self-supervised monocular scene flow estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7396–7405 (2020)
- Hur, J., Roth, S.: Self-supervised multi-frame monocular scene flow. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2684–2694 (2021)
- Ilg, E., Saikia, T., Keuper, M., Brox, T.: Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 614–630 (2018)
- Jaimez, M., Souiai, M., Gonzalez-Jimenez, J., Cremers, D.: A primal-dual framework for real-time dense rgb-d scene flow. In: 2015 IEEE international conference on robotics and automation (ICRA). pp. 98–104. IEEE (2015)
- Jampani, V., Kiefel, M., Gehler, P.V.: Learning sparse high dimensional filters: Image filtering, dense crfs and bilateral neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4452–4461 (2016)
- Jonschkowski, R., Stone, A., Barron, J.T., Gordon, A., Konolige, K., Angelova, A.: What matters in unsupervised optical flow. In: European Conference on Computer Vision. pp. 557– 572. Springer (2020)
- Kiefel, M., Jampani, V., Gehler, P.: Permutohedral lattice cnns. In: ICLR Workshop Track 2015 (2015)
- 15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Kittenplon, Y., Eldar, Y.C., Raviv, D.: Flowstep3d: Model unrolling for self-supervised scene flow estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4114–4123 (2021)
- Li, R., Lin, G., He, T., Liu, F., Shen, C.: Hcrf-flow: Scene flow from point clouds with continuous high-order crfs and position-aware flow embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 364–373 (2021)
- Li, R., Lin, G., Xie, L.: Self-point-flow: Self-supervised scene flow estimation from point clouds with optimal transport and random walk. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 15577–15586 (2021)

- 16 G. Wang et al.
- Liu, X., Qi, C.R., Guibas, L.J.: Flownet3d: Learning scene flow in 3d point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 529–537 (2019)
- Liu, X., Yan, M., Bohg, J.: Meteornet: Deep learning on dynamic 3d point cloud sequences. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9246– 9255 (2019)
- Liu, Z., Zhou, S., Suo, C., Yin, P., Chen, W., Wang, H., Li, H., Liu, Y.H.: Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2831–2840 (2019)
- Ma, W.C., Wang, S., Hu, R., Xiong, Y., Urtasun, R.: Deep rigid instance scene flow. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3614–3622 (2019)
- Mayer, N., Ilg, E., Hausser, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: 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. pp. 4040–4048 (2016)
- Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3061–3070 (2015)
- Menze, M., Heipke, C., Geiger, A.: Object scene flow. ISPRS Journal of Photogrammetry and Remote Sensing 140, 60–76 (2018)
- Mittal, H., Okorn, B., Held, D.: Just go with the flow: Self-supervised scene flow estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11177–11185 (2020)
- Pons, J.P., Keriven, R., Faugeras, O.: Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. International Journal of Computer Vision 72(2), 179–193 (2007)
- Puy, G., Boulch, A., Marlet, R.: Flot: Scene flow on point clouds guided by optimal transport. In: European conference on computer vision. pp. 527–544. Springer (2020)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
- 30. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. Advances in neural information processing systems **30** (2017)
- Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J.: Splatnet: Sparse lattice networks for point cloud processing. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2530–2539 (2018)
- Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018)
- Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: European conference on computer vision. pp. 402–419. Springer (2020)
- Ushani, A.K., Wolcott, R.W., Walls, J.M., Eustice, R.M.: A learning approach for real-time temporal scene flow estimation from lidar data. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 5666–5673. IEEE (2017)
- Valgaerts, L., Bruhn, A., Zimmer, H., Weickert, J., Stoll, C., Theobalt, C.: Joint estimation of motion, structure and geometry from stereo sequences. In: European Conference on Computer Vision. pp. 568–581. Springer (2010)
- Vedula, S., Baker, S., Rander, P., Collins, R., Kanade, T.: Three-dimensional scene flow. In: Proceedings of the Seventh IEEE International Conference on Computer Vision. vol. 2, pp. 722–729. IEEE (1999)

- Vogel, C., Schindler, K., Roth, S.: Piecewise rigid scene flow. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1377–1384 (2013)
- Vogel, C., Schindler, K., Roth, S.: 3d scene flow estimation with a piecewise rigid scene model. International Journal of Computer Vision 115(1), 1–28 (2015)
- Wang, G., Jiang, C., Shen, Z., Miao, Y., Wang, H.: Sfgan: Unsupervised generative adversarial learning of 3d scene flow from the 3d scene self. Advanced Intelligent Systems 4(4), 2100197 (2022)
- Wang, G., Liu, H., Chen, M., Yang, Y., Liu, Z., Wang, H.: Anchor-based spatio-temporal attention 3-d convolutional networks for dynamic 3-d point cloud sequences. IEEE Transactions on Instrumentation and Measurement **70**, 1–11 (2021)
- Wang, G., Peng, C., Zhang, J., Wang, H.: Interactive multi-scale fusion of 2d and 3d features for multi-object tracking. arXiv preprint arXiv:2203.16268 (2022)
- Wang, G., Tian, X., Ding, R., Wang, H.: Unsupervised learning of 3d scene flow from monocular camera. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 4325–4331. IEEE (2021)
- Wang, G., Wu, X., Jiang, S., Liu, Z., Wang, H.: Efficient 3d deep lidar odometry. arXiv preprint arXiv:2111.02135 (2021)
- Wang, G., Wu, X., Liu, Z., Wang, H.: Hierarchical attention learning of scene flow in 3d point clouds. IEEE Transactions on Image Processing 30, 5168–5181 (2021)
- Wang, G., Wu, X., Liu, Z., Wang, H.: Pwclo-net: Deep lidar odometry in 3d point clouds using hierarchical embedding mask optimization. In: Proc. CVPR. pp. 15910–15919 (2021)
- Wang, G., Yang, Y., Zhang, H., Liu, Z., Wang, H.: Spherical interpolated convolutional network with distance-feature density for 3-d semantic segmentation of point clouds. IEEE Transactions on Cybernetics (2021)
- 47. Wang, H., Pang, J., Lodhi, M.A., Tian, Y., Tian, D.: Festa: Flow estimation via spatialtemporal attention for scene point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14173–14182 (2021)
- Wang, S., Sun, Y., Liu, C., Liu, M.: Pointtracknet: An end-to-end network for 3-d object detection and tracking from point clouds. IEEE Robotics and Automation Letters 5(2), 3206– 3212 (2020)
- Wei, Y., Wang, Z., Rao, Y., Lu, J., Zhou, J.: 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. pp. 6954–6963 (2021)
- Wu, W., Wang, Z.Y., Li, Z., Liu, W., Fuxin, L.: Pointpwc-net: Cost volume on point clouds for (self-) supervised scene flow estimation. In: European conference on computer vision. pp. 88–107. Springer (2020)
- Xia, Y., Xu, Y., Li, S., Wang, R., Du, J., Cremers, D., Stilla, U.: Soe-net: A self-attention and orientation encoding network for point cloud based place recognition. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition. pp. 11348–11357 (2021)
- Yang, G., Ramanan, D.: Upgrading optical flow to 3d scene flow through optical expansion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1334–1343 (2020)
- Yang, G., Ramanan, D.: Learning to segment rigid motions from two frames. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1266–1275 (2021)