

PV-RAFT: Point-Voxel Correlation Fields for Scene Flow Estimation of Point Clouds

Yi Wei^{1,2,3,*}, Ziyi Wang^{1,2,3,*}, Yongming Rao^{1,2,3,*}, Jiwen Lu^{1,2,3,†}, Jie Zhou^{1,2,3,4}

¹Department of Automation, Tsinghua University, China

²State Key Lab of Intelligent Technologies and Systems, China

³Beijing National Research Center for Information Science and Technology, China

⁴Tsinghua Shenzhen International Graduate School, Tsinghua University, China

{y-wei19, wzyi20}@mails.tsinghua.edu.cn; raoyongming95@gmail.com; {lujiwen, jzhou}@tsinghua.edu.cn

Abstract

In this paper, we propose a Point-Voxel Recurrent All-Pairs Field Transforms (PV-RAFT) method to estimate scene flow from point clouds. Since point clouds are irregular and unordered, it is challenging to efficiently extract features from all-pairs fields in the 3D space, where all-pairs correlations play important roles in scene flow estimation. To tackle this problem, we present point-voxel correlation fields, which capture both local and long-range dependencies of point pairs. To capture point-based correlations, we adopt the K-Nearest Neighbors search that preserves fine-grained information in the local region. By voxelizing point clouds in a multi-scale manner, we construct pyramid correlation voxels to model long-range correspondences. Integrating these two types of correlations, our PV-RAFT makes use of all-pairs relations to handle both small and large displacements. We evaluate the proposed method on the FlyingThings3D and KITTI Scene Flow 2015 datasets. Experimental results show that PV-RAFT outperforms state-of-the-art methods by remarkable margins.

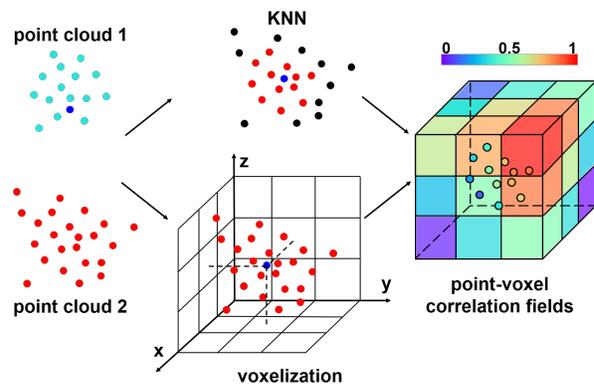


Figure 1: Illustration of the proposed point-voxel correlation fields. For a point in the source point cloud, we find its k -nearest neighbors in the target point cloud to extract point-based correlations. Moreover, we model long-range interactions by building voxels centered around this source point. Combining these two types of correlations, our PV-RAFT captures all-pairs dependencies to deal with both large and small displacements.

1. Introduction

3D scene understanding [7, 16, 32, 34, 45, 51] has attracted more and more attention in recent years due to its wide real-world applications. As one fundamental 3D computer vision task, scene flow estimation [6, 10, 18, 23, 26, 50] focuses on computing the 3D motion field between two consecutive frames, which provides important dynamic information. Conventionally, scene flow is directly estimated from RGB images [20, 21, 41, 43]. Since 3D data becomes easier to obtain, many works [6, 18, 26, 50] begin to focus on scene flow estimation of point clouds more recently.

*Equal Contribution

†Corresponding author

Thanks to the recent advances in deep learning, many approaches adopt deep neural networks for scene flow estimation [6, 18, 26, 39, 50]. Among these methods, [18, 50] borrow ideas from [5, 11, 35], leveraging techniques in mature optical flow area. FlowNet3D designs a flow embedding module to calculate correlations between two frames. Built upon PWC-Net [35], PointPWC-Net [50] introduces a learnable point-based cost volume without the need of 4D dense tensors. These methods follow a coarse-to-fine strategy, where scene flow is first computed at low resolution and then upsampled to high resolution. However, this strategy has several limitations [37], e.g. error accumulation from early steps and the tendency to miss fast-moving objects. One possible solution is to adopt Recurrent All-Pairs

Field Transforms (RAFT) [37], a state-of-the-art method for 2D optical flow, that builds correlation volumes for all pairs of pixels. Compared with the coarse-to-fine strategy, the all-pairs field preserves both local correlations and long-range relations. Nevertheless, it is non-trivial to lift it to the 3D space. Due to the irregularity of point clouds, building structured all-pairs correlation fields becomes challenging. Moreover, since point clouds are unordered, it is difficult to efficiently look up neighboring points of a 3D position. Unfortunately, the correlation volumes used in previous methods [6, 18, 50] only consider near neighbors, which fails to capture all-pairs relations.

To address these issues, we present point-voxel correlation fields that aggregate the advantages of both point-based and voxel-based correlations (illustrated in Figure 1). As mentioned in [19, 32, 36], point-based features maintain fine-grained information while voxel-based operation efficiently encodes large point set. Motivated by this fact, we adopt K-Nearest Neighbor (KNN) search to find a fixed number of neighboring points for point-based correlation fields. Meanwhile, we voxelize target point clouds in a multi-scale fashion to build pyramid correlation voxels. These voxel-based correlation fields collect long-term dependencies and guide the predicted direction. Moreover, to save memory, we present a truncation mechanism to abandon the correlations with low scores.

Based on point-voxel correlation fields, we propose a Point-Voxel Recurrent All-Pairs Field Transforms (PV-RAFT) method to construct a new network architecture for scene flow estimation of point clouds. Our method first employs a feature encoder to extract per-point features, which are utilized to build all-pair correlation fields. Then we adopt a GRU-based operator to update scene flow in an iterative manner, where we leverage both point-based and voxel-based mechanisms to look up correlation features. Finally, a refinement module is introduced to smooth the estimated scene flow. To evaluate our method, we conducted extensive experiments on the FlyingThings3D [20] and KITTI [21, 22] datasets. Results show that our PV-RAFT outperforms state-of-the-art methods by a large margin. The code is available at <https://github.com/weiyithu/PV-RAFT>.

2. Related Work

3D Deep Learning: Increased attention has been paid to 3D deep learning [7, 12, 16, 27–29, 31–33, 45, 49, 51] due to its wide applications. As a pioneer work, PointNet [28] is the first deep learning framework directly operating on point clouds. It uses a max pooling layer to aggregate features of unordered set. PointNet++ [29] introduces a hierarchical structure by using PointNet as a unit module. Kd-network [14] equips a kd-tree to divide point clouds and compute a sequence of hierarchical representations.

DGCNN [46] models point clouds as a graph and utilizes graph neural networks to extract features. Thanks to these architectures, great achievements have been made in many 3D areas, *e.g.* 3D recognition [15, 17, 28, 29], 3D segmentation [7, 12, 45]. Recently, several works [19, 32, 36] simultaneously leverage point-based and voxel-based methods to operate on point clouds. Liu *et al.* [19] present Point-Voxel CNN (PVCNN) for efficient 3D deep learning. It combines voxel-based CNN and point-based MLP to extract features. As a follow-up, Tang *et al.* [36] design SPVConv [36] which adopts Sparse Convolution with the high-resolution point-based network. They further propose 3D-NAS to search the best architecture. PV-RCNN [32] takes advantage of high-quality 3D proposals from 3D voxel CNN and accurate location information from PointNet-based set abstraction operation. Instead of equipping point-voxel architecture to extract features, we design point-voxel correlation fields to capture correlations.

Optical Flow Estimation: Optical flow estimation [5, 9, 11, 30, 38] is a hot topic in 2D area. FlowNet [5] is the first trainable CNN for optical flow estimation, adopting a U-Net autoencoder architecture. Based on [5], FlowNet2 [11] stacks several FlowNet models to compute large-displacement optical flows. With this cascaded backbone, FlowNet2 [11] outperforms FlowNet [5] by a large margin. To deal with large motions, SPyNet [30] adopts the coarse-to-fine strategy with a spatial pyramid. Beyond SPyNet [30], PWC-Net [35] builds a cost volume by limiting the search range at each pyramid level. Similar to PWC-Net, LiteFlowNet [9] also utilizes multiple correlation layers operating on a feature pyramid. Recently, GLU-Net [38] combines global and local correlation layers with an adaptive resolution strategy, which achieves both high accuracy and robustness. Different from the coarse-to-fine strategy, RAFT [37] constructs the multi-scale 4D correlation volume for all pairs of pixels. It further updates the flow field through a recurrent unit iteratively, and achieves state-of-the-art performance on optical flow estimation task. The basic structure of our PV-RAFT is similar to theirs. However, we adjust the framework to fit point clouds data format and propose point-voxel correlation fields to leverage all-pairs relations.

Scene Flow Estimation: First introduced in [41], scene flow is the three-dimension vector to describe the motion in real scenes. Beyond this pioneer work, many studies estimate scene flow from RGB images [1, 2, 8, 25, 42–44, 47, 48]. Based on stereo sequences, [8] proposes a variational method to estimate scene flow. Similar to [8], [48] decouples the position and velocity estimation steps with consistent displacements in the stereo images. [44] represents dynamic scenes as a collection of rigidly moving planes and accordingly introduces a piecewise rigid scene model. With the development of 3D sensors, it becomes

easier to get high-quality 3D data. More and more works focus on how to leverage point clouds for scene flow estimation [4, 6, 18, 26, 39, 40, 50]. FlowNet3D [18] introduces two layers to simultaneously learn deep hierarchical features of point clouds and flow embeddings. Inspired by Bilateral Convolutional Layers, HPLFlowNet [6] projects unstructured point clouds onto a permutohedral lattice. Operating on permutohedral lattice points, it can efficiently calculate scene flow. Benefiting from the coarse-to-fine strategy, PointPWC-Net [50] proposes cost volume, upsampling, and warping layers for scene flow estimation. Different from the above methods, FLOT [26] adopts the optimal transport to find correspondences. However, the correlation layers introduced in these methods only consider the neighbors in a local region, which fail to efficiently capture long-term dependencies. With point-voxel correlation fields, our PV-RAFT captures both local and long-range correlations.

3. Approach

To build all-pairs fields, it is important to design a correlation volume which can capture both short-range and long-range relations. In this section, we first explain how to construct point-voxel correlation fields on point clouds. Then we will introduce the pipeline of our Point-Voxel Recurrent All-Pairs Field Transforms (PV-RAFT).

3.1. Point-Voxel Correlation Fields

We first construct a full correlation volume based on feature similarities between all pairs. Given point clouds features $E_\theta(P_1) \in \mathbb{R}^{N_1 \times D}$, $E_\theta(P_2) \in \mathbb{R}^{N_2 \times D}$, where D is the feature dimension, the correlation fields $\mathbf{C} \in \mathbb{R}^{N_1 \times N_2}$ can be easily calculated by matrix dot product:

$$\mathbf{C} = E_\theta(P_1) \cdot E_\theta(P_2) \quad (1)$$

Correlation Lookup: The correlation volume \mathbf{C} is built only once and is kept as a lookup table for flow estimations in different steps. Given a source point $p_1 = (x_1, y_1, z_1) \in P_1$, a target point $p_2 = (x_2, y_2, z_2) \in P_2$ and an estimated scene flow $f = (f_1, f_2, f_3) \in \mathbf{f}$, the source point is expected to move to $q = (x_1 + f_1, x_2 + f_2, x_3 + f_3) \in Q$, where Q is the translated point cloud. We can easily get the correlation fields between Q and P_2 by searching the neighbors of Q in P_2 and looking up the corresponding correlation values in \mathbf{C} . Such looking-up procedure avoids extracting features of Q and calculating matrix dot product repeatedly while keeping the all-pairs correlations available at the same time. Since 3D points data is not structured in the dense voxel, grid sampling is no longer useful and we cannot directly convert 2D method [37] into 3D version. Thus, the main challenge is how to locate neighbors and look up correlation values efficiently in the 3D space.

Truncated Correlation: According to our experimental results, not all correlation entries are useful in the subsequent

correlation lookup process. The pairs with higher similarity often guide the correct direction of flow estimation, while dissimilar pairs tend to make little contribution. To save memory and increase calculation efficiency in correlation lookup, for each point in P_1 , we select its top- M highest correlations. Specifically, we will get truncated correlation fields $\mathbf{C}_M \in \mathbb{R}^{N_1 \times M}$, where $M < N_2$ is the pre-defined truncation number. The point branch and voxel branch are built upon truncated correlation fields.

Point Branch: A common practice to locate neighbors in 3D point clouds is to use K-Nearest Neighbors (KNN) algorithm. Suppose the top- k nearest neighbors of Q in P_2 is $\mathcal{N}_k = \mathcal{N}(Q)_k$ and their corresponding correlation values are $\mathbf{C}_M(\mathcal{N}_k)$, the correlation feature between Q and P_2 can be defined as:

$$\mathbf{C}_p(Q, P_2) = \max_k (\text{MLP}(\text{concat}(\mathbf{C}_M(\mathcal{N}_k), \mathcal{N}_k - Q))) \quad (2)$$

where *concat* stands for concatenation and *max* indicates a max pooling operation on k dimension. We briefly note $\mathcal{N}(Q)$ as \mathcal{N} in the following statements as all neighbors are based on Q in this paper. The point branch extracts fine-grained correlation features of the estimated flow since the nearest neighbors are often close to the query point, illustrated in the upper branch of Figure 1. While the point branch is able to capture local correlations, long-range relations are often not taken into account in KNN scenario. Existing methods try to solve this problem by implementing the coarse-to-fine strategy, but error often accumulates if estimates in the coarse stage are not accurate.

Voxel Branch: To tackle the problem mentioned above, we propose a voxel branch to capture long-range correlation features. Instead of voxelizing Q directly, we build voxel neighbor cubes centered around Q and check which points in P_2 lie in these cubes. Moreover, we also need to know each point’s relative direction to Q . Therefore, if we denote sub-cube side length by r and cube resolution by a , then the neighbor cube of Q would be a $a \times a \times a$ Rubik’s cube:

$$\mathcal{N}_{r,a} = \{\mathcal{N}_r^{(i)} | \mathbf{i} \in \mathbb{Z}^3\} \quad (3)$$

$$\mathcal{N}_r^{(i)} = \{Q + \mathbf{i} * r + \mathbf{dr} | \|\mathbf{dr}\|_1 \leq \frac{r}{2}\} \quad (4)$$

where $\mathbf{i} = [i, j, k]^T$, $[-\frac{a}{2}] \leq i, j, k \leq [\frac{a}{2}] \in \mathbb{Z}$ and each $r \times r \times r$ sub-cube $\mathcal{N}_r^{(i)}$ indicates a specific direction of neighbor points (e.g., $[0, 0, 0]^T$ indicates the central sub-cube). Then we identify all neighbor points in the sub-cube $\mathcal{N}_r^{(i)}$ and average their correlation values to get sub-cube features. The correlation feature between Q and P_2 can be defined as:

$$\mathbf{C}_v(Q, P_2) = \text{MLP} \left(\text{concat}_{\mathbf{i}} \left(\frac{1}{n_{\mathbf{i}}} \sum_{n_{\mathbf{i}}} \mathbf{C}_M(\mathcal{N}_r^{(i)}) \right) \right) \quad (5)$$

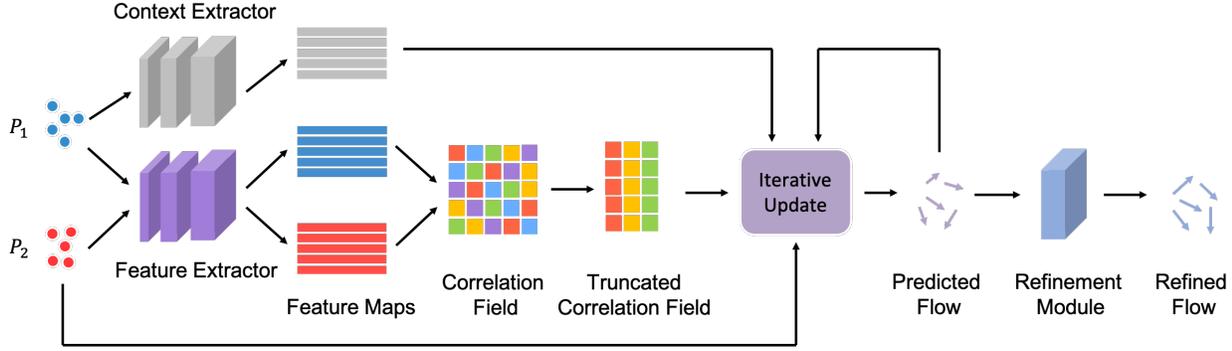


Figure 2: Illustration of the proposed PV-RAFT architecture. The feature extractor encodes high dimensional features of both P_1 and P_2 , while the context extractor only encodes context features of P_1 . We calculate the matrix dot product of two feature maps to construct all-pair correlation fields. The truncated correlation field is then used in iterative update block to save memory. The detailed structure of 'Iterative Update' module can be found in Figure 3. The predicted flow from the iteration block finally converges to a static status and is fed into the separately trained refinement module. We use the refined flow as the final scene flow prediction.

where n_i is the number of points in P_2 that lie in the i^{th} sub-cube of Q and $\mathbf{C}_v(Q, P_2) \in \mathbb{R}^{N_1 \times a^3}$. Please refer to the lower branch of Figure 1 for illustration.

The Voxel branch helps to capture long-range correlation features as r, a could be large enough to cover distant points. Moreover, we propose to extract pyramid correlation voxels with fixed cube resolution a and proportionate growing sub-cube side length r . During each pyramid iteration, r is doubled so that the neighbor cube expands to include farther points. The pyramid features are concatenated together before feeding into the MLP layer.

3.2. PV-RAFT

Given the proposed correlation fields that combine the fine-grained and long-range features, we build a deep neural network for scene flow estimation. The pipeline consists of four stages: (1) feature extraction, (2) correlation fields construction, (3) iterative scene flow estimation, (4) flow refinement. The first three stages are differentiable in an end-to-end manner, while the fourth one is trained separately with previous parts frozen. Our framework is called PV-RAFT and in this section we will introduce it in detail. Please refer to Figure 2 for illustration.

Feature Extraction: The feature extractor E_θ encodes point clouds with mere coordinates information into higher dimensional feature space, as $E_\theta : \mathbb{R}^{n \times 3} \mapsto \mathbb{R}^{n \times D}$. Our backbone framework is based on PointNet++ [29]. For consecutive point clouds input P_1, P_2 , the feature extractor outputs $E_\theta(P_1), E_\theta(P_2)$ as backbone features. Besides, we design a content feature extractor E_γ to encode context feature of P_1 . Its structure is exactly the same as feature extractor E_θ , without weight sharing. The output context feature $E_\gamma(P_1)$ is used as auxiliary context information in GRU iteration.

Correlation Fields Construction: As is introduced in Section 3.1, we build all-pair correlation fields \mathbf{C} based on backbone features $E_\theta(P_1), E_\theta(P_2)$. Then we truncate it according to correlation value sorting and keep it as a lookup table for later iterative updates.

Iterative Flow Estimation: The iterative flow estimation begins with the initialize state $\mathbf{f}_0 = 0$. With each iteration, the scene flow estimation is updated upon the current state: $\mathbf{f}_{t+1} = \mathbf{f}_t + \Delta \mathbf{f}$. Eventually, the sequence converges to the final prediction $\mathbf{f}_T \rightarrow \mathbf{f}^*$. Each iteration takes the following variables as input: (1) correlation features, (2) current flow estimate, (3) hidden states from the previous iteration, (4) context features. First, the correlation features are the combination of both fine-grained point-based ones and long-range pyramid-voxel-based ones:

$$\mathbf{C}_t = \mathbf{C}_p(Q_t, P_2) + \mathbf{C}_v(Q_t, P_2) \quad (6)$$

Second, the current flow estimation is simply the direction vector between Q_t and P_1 :

$$\mathbf{f}_t = Q_t - P_1 \quad (7)$$

Third, the hidden state h_t is calculated by GRU cell [37]:

$$z_t = \sigma(\text{Conv}_{1d}([h_{t-1}, x_t], W_z)) \quad (8)$$

$$r_t = \sigma(\text{Conv}_{1d}([h_{t-1}, x_t], W_r)) \quad (9)$$

$$\hat{h}_t = \tanh(\text{Conv}_{1d}([r_t \odot h_{t-1}, x_t], W_h)) \quad (10)$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \hat{h}_t \quad (11)$$

where x_t is a concatenation of correlation \mathbf{C}_t , current flow \mathbf{f}_t and context features $E_\gamma(P_1)$. Finally, the hidden state h_t is fed into a small convolutional network to get the final scene flow estimate \mathbf{f}^* . The detailed iterative update process is illustrated in Figure 3.

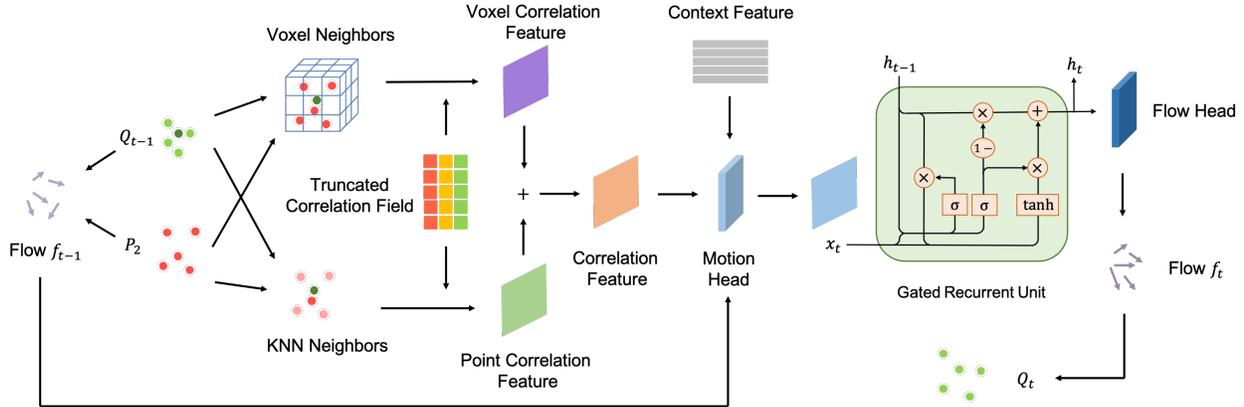


Figure 3: Illustration of the iterative update. This figure is a detailed explanation of the ‘Iterative Update’ module in Figure 2. During iteration t , we find both voxel neighbors and KNN of Q_{t-1} in P_2 . This helps us extract long-range voxel correlation features and fine-grained point correlation features from the truncated correlation field. The combined correlation feature, together with context feature and current flow estimate f_{t-1} are fed to a convolutional motion head. The output is used as x_t of the Gated Recurrent Unit (GRU). Finally, the flow head encodes the hidden state h_t of GRU to predict the residual of flow estimation, which is used to update f_t and Q_t .

Flow Refinement: The purpose of designing this flow refinement module is to make scene flow prediction \mathbf{f}^* smoother in the 3D space. Specifically, the estimated scene flow from previous stages is fed into three convolutional layers and one fully connected layer. To update flow for more iterations without out of memory, the refinement module is not trained end-to-end with other modules. We first train the backbone and iterative update module, then we freeze the weights and train the refinement module alone.

3.3. Loss Function

Flow Supervision: We follow the common practice of supervised scene flow learning to design our loss function. In detail, we use l_1 -norm between the ground truth flow \mathbf{f}_{gt} and estimated flow \mathbf{f}_{est} for each iteration:

$$\mathcal{L}_{iter} = \sum_{t=1}^T w^{(t)} \|\mathbf{f}_{est}^{(t)} - \mathbf{f}_{gt}\|_1 \quad (12)$$

where T is the total amount of iterative updates, $\mathbf{f}_{est}^{(t)}$ is the flow estimate at t^{th} iteration, and $w^{(t)}$ is the weight of t^{th} iteration:

$$w^{(t)} = \gamma * (T - t - 1) \quad (13)$$

where γ is a hyper-parameter and we set $\gamma = 0.8$ in our experiments.

Refinement Supervision: When we freeze the weights of previous stages and only train the refinement module, we design a similar refinement loss:

$$\mathcal{L}_{ref} = \|\mathbf{f}_{ref} - \mathbf{f}_{gt}\|_1 \quad (14)$$

where \mathbf{f}_{ref} is the flow prediction from refinement module.

4. Experiments

In this section, we conducted extensive experiments to verify the superiority of our PV-RAFT. We first introduce the experimental setup, including datasets, implementation details and evaluation metrics. Then we show main results on the FlyingThings3D [20] and KITTI [21, 22] datasets, as well as ablation studies. Finally, we give a further analysis of PV-RAFT to better illustrate the effectiveness of our proposed method.

4.1. Experimental Setup

Datasets: Same with [6, 18, 26, 50], we trained our model on the FlyingThings3D [20] dataset and evaluated it on both FlyingThings3D [20] and KITTI [21, 22] datasets. We followed [6] to preprocess data. As a large-scale synthetic dataset, FlyingThings3D is the first benchmark for scene flow estimation. With the objects from ShapeNet [3], FlyingThings3D consists of rendered stereo and RGB-D images. Totally, there are 19,640 pairs of samples in the training set and 3,824 pairs in the test set. Besides, we kept aside 2000 samples from the training set for validation. We lifted depth images to point clouds and optical flow to scene flow instead of operating on RGB images. As another benchmark, KITTI Scene Flow 2015 is a dataset for scene flow estimation in real scans [21, 22]. It is built from KITTI raw data by annotating dynamic motions. Following previous works [6, 18, 26, 50], we evaluated on 142 samples in the training set since point clouds were not available in the test set. Ground points were removed by height (0.3m). Further, we deleted points whose depths are larger than 35m.

Implementation Details: We randomly sampled 8192 points in each point cloud to train PV-RAFT. For the point

Table 1: Performance comparison on the FlyingThings3D and KITTI datasets. All methods are trained on FlyingThings3D in a supervised manner. The best results for each dataset are marked in bold.

Dataset	Method	EPE(m)↓	Acc Strict↑	Acc Relax↑	Outliers↓
FlyingThings3D	FlowNet3D [18]	0.1136	0.4125	0.7706	0.6016
	HPLFlowNet [6]	0.0804	0.6144	0.8555	0.4287
	PointPWC-Net [50]	0.0588	0.7379	0.9276	0.3424
	FLOT [26]	0.052	0.732	0.927	0.357
	PV-RAFT	0.0461	0.8169	0.9574	0.2924
KITTI	FlowNet3D [18]	0.1767	0.3738	0.6677	0.5271
	HPLFlowNet [6]	0.1169	0.4783	0.7776	0.4103
	PointPWC-Net [50]	0.0694	0.7281	0.8884	0.2648
	FLOT [26]	0.056	0.755	0.908	0.242
	PV-RAFT	0.0560	0.8226	0.9372	0.2163

branch, we searched 32 nearest neighbors. For the voxel branch, we set cube resolution $a = 3$ and built 3-level pyramid with $r = 0.25, 0.5, 1$. To save memory, we set truncation number M as 512. We updated scene flow for 8 iterations during training and evaluated the model with 32 flow updates. The backbone and iterative module were trained for 20 epochs. Then, we fixed their weights with 32 iterations and trained the refinement module for another 10 epochs. PV-RAFT was implemented in PyTorch [24]. We utilized Adam optimizer [13] with initial learning rate as 0.001.

Evaluation Metrics: We adopted four evaluation metrics used in [6, 18, 26, 50], including EPE, Acc Strict, Acc Relax and Outliers. We denote estimated scene flow and ground-truth scene flow as f_{est} and f_{gt} respectively. The evaluation metrics are defined as follows:

- **EPE:** $\|f_{est} - f_{gt}\|_2$. The end point error averaged on each point in meters.
- **Acc Strict:** the percentage of points whose **EPE** $< 0.05m$ or relative error $< 5\%$.
- **Acc Relax:** the percentage of points whose **EPE** $< 0.1m$ or relative error $< 10\%$.
- **Outliers:** the percentage of points whose **EPE** $> 0.3m$ or relative error $> 10\%$.

4.2. Main Results

Quantitative results on the FlyingThings3D and KITTI datasets are shown in Table 1. Our PV-RAFT achieves state-of-the-art performances on both datasets, which verifies its superiority and generalization ability. Especially, for Outliers metric, our method outperforms FLOT by 18.1% and 10.6% on two datasets respectively. The qualitative results in Figure 4 further demonstrate the effectiveness of PV-RAFT. The first row and second row present visualizations on the FlyingThings3D and KITTI datasets respectively. As we can see, benefiting from point-voxel correlation fields, our method can accurately predict both small and large displacements.

4.3. Ablation Studies

We conducted experiments to confirm the effectiveness of each module in our method. Point-based correlation, voxel-based correlation and refinement module were applied to our framework incrementally. From Table 2, we can conclude that each module plays an important part in the whole pipeline. As two baselines, the methods with only point-based correlation or voxel-based correlation fail to achieve high performance, since they cannot capture all-pairs relations. An intuitive solution is to employ more nearest neighbors in the point branch to increase the receptive field or decrease the side length r in the voxel branch to take fine-grained correlations. However, we find that such straightforward methods lead to inferior results (See details in the supplemental material).

To better illustrate the effects of two types of correlations, we show visualizations in Figure 5. At the beginning of update steps, when predicted flows are initialized as zero, the estimated translated points are far from ground-truth correspondences in the target point cloud (first column). Under this circumstance, the similarity scores with near neighbors are small, where point-based correlation provides invalid information. In contrast, since voxel-based correlation has the large receptive field, it is able to find long-range correspondences and guide the prediction direction. As the update iteration increases, we will get more and more accurate scene flow. When translated points are near to the ground-truth correspondences, high-score correlations will concentrate on the centered lattice of the voxel (third column), which does not serve detailed correlations. However, we will get informative correlations from the point branch since KNN perfectly encodes local information.

4.4. Further Analysis

Effects of Truncation Operation: We introduce the truncation operation to reduce running memory while maintain the performance. To prove this viewpoint, we conducted

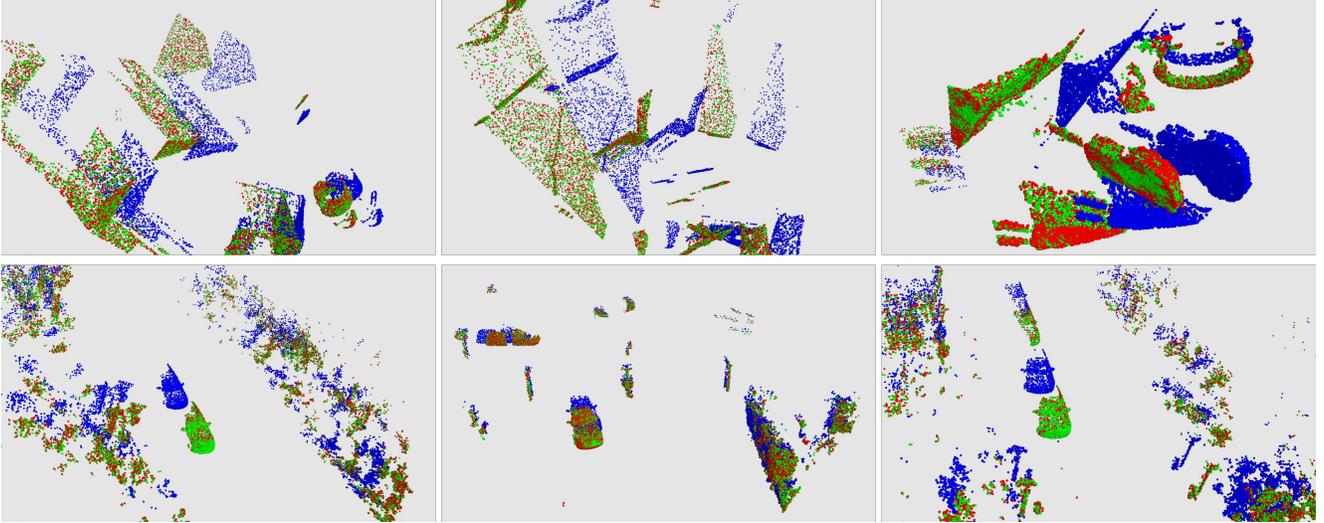


Figure 4: Qualitative results on FlyingThings3D (top) and KITTI (bottom). Blue points and red points indicate P_1 and P_2 respectively. Translated points $P_1 + \mathbf{f}$ are in green. Our PV-RAFT can deal with both small and large displacements’ cases.

Table 2: Ablation Studies of PV-RAFT on the FlyingThings3D dataset. We incrementally applied point-based correlation, voxel-based correlation and refinement module to the framework.

point-based correlation	voxel-based correlation	refine module	EPE(m)↓	Acc Strict↑	Acc Relax↑	Outliers↓
✓			0.0741	0.6111	0.8868	0.4549
	✓		0.0712	0.6146	0.8983	0.4492
✓	✓		0.0534	0.7348	0.9418	0.3645
✓	✓	✓	0.0461	0.8169	0.9574	0.2924

Table 3: Effects of truncation operation. M denotes the truncation number.

M	memory	EPE(m)↓	Acc Strict↑	Outliers↓
128	7.4G	0.0585	0.7113	0.3810
512	10.7G	0.0461	0.8169	0.2924
1024	14.1G	0.0475	0.8173	0.2910

Table 4: Comparison with other correlation volume methods. ”MLP+Maxpool” and ”patch-to-patch” are correlation volumes used in FlowNet3D [18] and PointPWC-Net [50] respectively.

Method	EPE(m)↓	Acc Strict↑	Outliers↓
MLP+Maxpool [18]	0.0704	0.7137	0.3843
patch-to-patch [50]	0.0614	0.7209	0.3628
point-voxel	0.0461	0.8169	0.2924

experiments with different truncation numbers M , which are shown in Table 3. On the one hand, when M is too small, the accuracy will degrade due to the lack of correlation information. On the other hand, achieving the comparable performance with $M = 512$, the model adopting $M = 1024$ needs about 14G running memory, which is not

available on many GPU services (e.g. RTX 2080 Ti). This result indicates that top 512 correlations are enough to accurately estimate scene flow with high efficiency.

Comparison with Other Correlation Volumes: To further demonstrate the superiority of the proposed point-voxel correlation fields, we did comparison with correlation volume methods introduced in FlowNet3D [18] and PointPWC-Net [50]. To fairly compare, we applied their correlation volumes in our framework to substitute point-voxel correlation fields. Evaluation results are shown in Table 4. Leveraging all-pairs relations, our point-voxel correlation module outperforms other correlation volume methods.

5. Conclusion

In this paper, we have proposed a PV-RAFT method for scene flow estimation of point clouds. With the point-voxel correlation fields, our method integrates two types of correlations and captures all-pairs relations. Leveraging the truncation operation and the refinement module, our framework becomes more accurate. Experimental results on the FlyingThings3D and KITTI datasets verify the superiority and generalization ability of PV-RAFT.

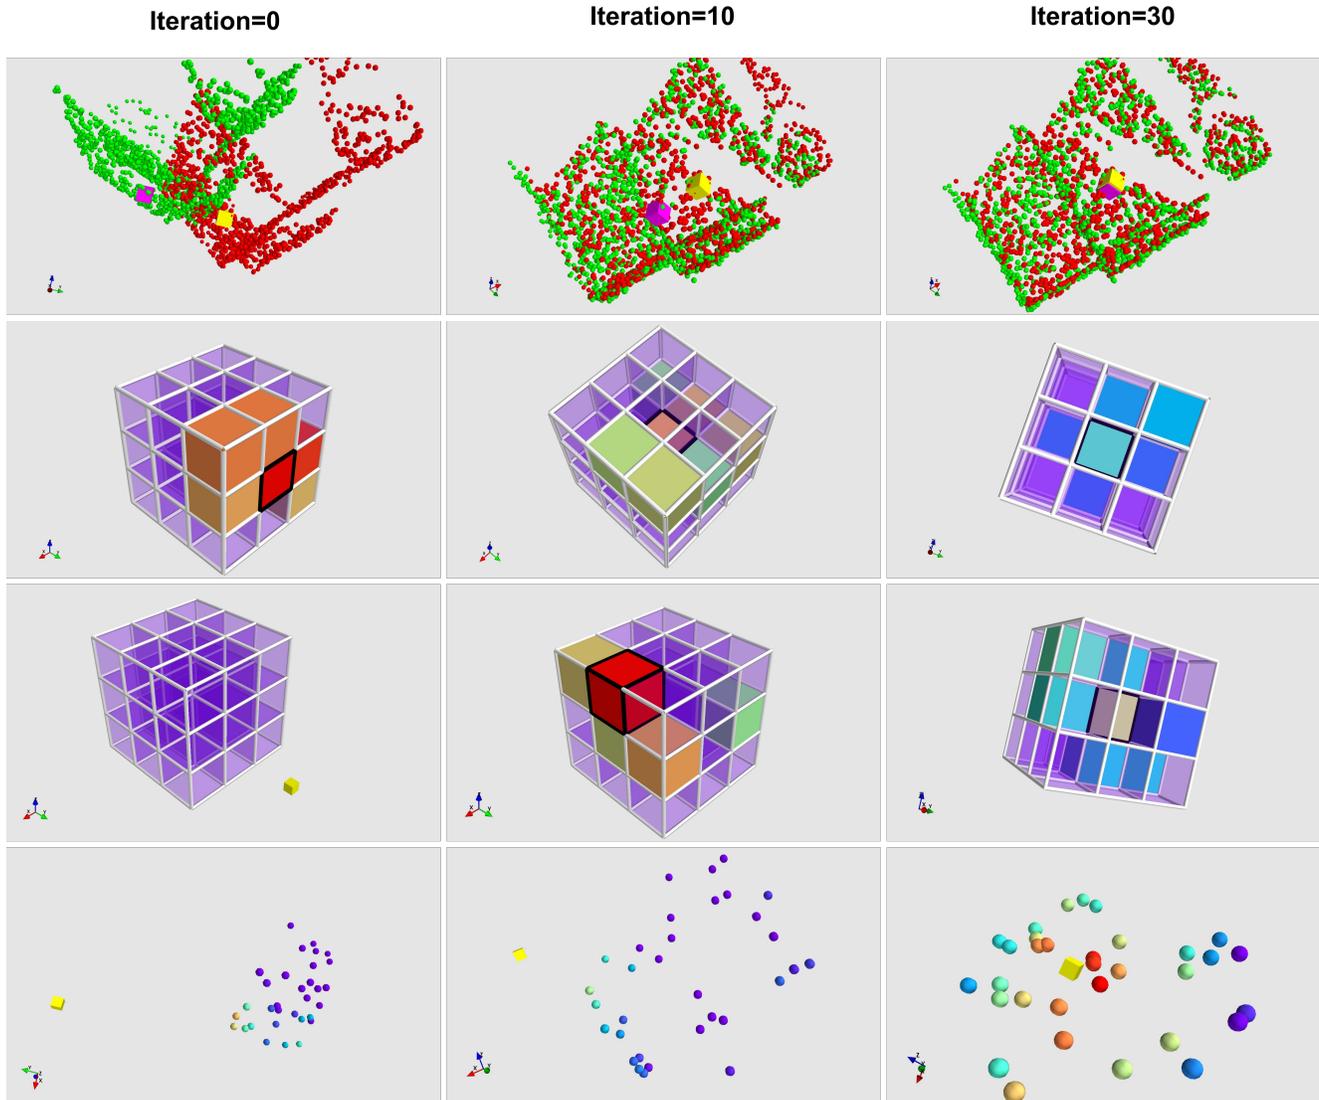


Figure 5: Visualization of point-voxel correlation fields. In the first row, green points represent translated point cloud $P_1 + \mathbf{f}$ while red points stand for target point cloud P_2 . The pink cube is a point in the translated point cloud, whose correspondence in P_2 is the yellow cube. The correlation fields of voxel branch are illustrated in the second ($r = 1$) and third ($r = 0.25$) rows. If the target point (yellow cube) lies in a lattice, the boundaries of this lattice will be colored in black. The last row exhibits the correlation field of the point branch. The colors of the last three rows indicate normalized correlation scores, where red is highest and purple is lowest (Figure 1 shows colormap). **At the beginning of the iterative update** (the first column), the predicted flow is not accurate so that the translated point is far from the target point. Since the voxel branch has large receptive fields, it can cover the target point while the point branch fails. From the first column and the second row, we see that the sub-cube which contains the target point has the highest correlation score. This indicates that the voxel branch provides effective guidance for flow prediction at early iterations. **As the iteration goes on**, the translated point gets near to the target point (the third column). The voxel branch only provides the coarse position of the target point (at the central sub-cube) while the point branch can accurately localize the target point by computing correlation scores of all neighbor points in the local region. The viewpoints are chosen to best visualize the sub-cube with the highest score.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant U1713214, Grant

U1813218, Grant 61822603, in part by Beijing Academy of Artificial Intelligence (BAAI), and in part by a grant from the Institute for Guo Qiang, Tsinghua University.

References

- [1] Tali Basha, Yael Moses, and Nahum Kiryati. Multi-view scene flow estimation: A view centered variational approach. *IJCV*, 2013. 2
- [2] Jan Čech, Jordi Sanchez-Riera, and Radu Horaud. Scene flow estimation by growing correspondence seeds. In *CVPR*, 2011. 2
- [3] 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
- [4] Ayush Dewan, Tim Caselitz, Gian Diego Tipaldi, and Wolfram Burgard. Rigid scene flow for 3d lidar scans. In *IROS*, 2016. 3
- [5] 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 *CVPR*, 2015. 1, 2
- [6] 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 *CVPR*, 2019. 1, 2, 3, 5, 6
- [7] Ji Hou, Angela Dai, and Matthias Nießner. 3d-sis: 3d semantic instance segmentation of rgb-d scans. In *CVPR*, 2019. 1, 2
- [8] Frédéric Huguet and Frédéric Devernay. A variational method for scene flow estimation from stereo sequences. In *ICCV*, 2007. 2
- [9] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Lite-flowNet: A lightweight convolutional neural network for optical flow estimation. In *CVPR*, 2018. 2
- [10] Junhwa Hur and Stefan Roth. Self-Supervised Monocular Scene Flow Estimation. In *CVPR*, 2020. 1
- [11] 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 *CVPR*, 2017. 1, 2
- [12] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. In *CVPR*, 2020. 2
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [14] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *ICCV*, 2017. 2
- [15] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *CVPR*, 2018. 2
- [16] Peiliang Li, Xiaozhi Chen, and Shaojie Shen. Stereo r-cnn based 3d object detection for autonomous driving. In *CVPR*, 2019. 1, 2
- [17] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *NeurIPS*, 2018. 2
- [18] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. FlowNet3d: Learning scene flow in 3d point clouds. In *CVPR*, 2019. 1, 2, 3, 5, 6, 7
- [19] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. In *NeurIPS*, 2019. 2
- [20] 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 *CVPR*, 2016. 1, 2, 5
- [21] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, 2015. 1, 2, 5
- [22] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3d estimation of vehicles and scene flow. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 2015. 2, 5
- [23] Himangi Mittal, Brian Okorn, and David Held. Just go with the flow: Self-supervised scene flow estimation. In *CVPR*, 2020. 1
- [24] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [25] Jean-Philippe Pons, Renaud Keriven, and Olivier Faugeras. Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. *IJCV*, 2007. 2
- [26] Gilles Puy, Alexandre Boulch, and Renaud Marlet. FLOT: Scene Flow on Point Clouds guided by Optimal Transport. In *ECCV*, 2020. 1, 3, 5, 6
- [27] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, 2019. 2
- [28] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 2
- [29] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017. 2, 4
- [30] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, 2017. 2
- [31] Yongming Rao, Jiwen Lu, and Jie Zhou. Global-local bidirectional reasoning for unsupervised representation learning of 3d point clouds. In *CVPR*, 2020. 2
- [32] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection. In *CVPR*, 2020. 1, 2
- [33] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3d object proposal generation and detection from point cloud. In *CVPR*, 2019. 2
- [34] Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. SplatNet: Sparse lattice networks for point cloud processing. In *CVPR*, 2018. 1
- [35] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 1, 2

- [36] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *ECCV*, 2020. [2](#)
- [37] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *ECCV*, 2020. [1](#), [2](#), [3](#), [4](#)
- [38] Prune Truong, Martin Danelljan, and Radu Timofte. GLU-Net: Global-Local Universal Network for Dense Flow and Correspondences. In *CVPR*, 2020. [2](#)
- [39] Arash K Ushani and Ryan M Eustice. Feature learning for scene flow estimation from lidar. In *Conference on Robot Learning*, 2018. [1](#), [3](#)
- [40] Arash K Ushani, Ryan W Wolcott, Jeffrey M Walls, and Ryan M Eustice. A learning approach for real-time temporal scene flow estimation from lidar data. In *ICRA*, 2017. [3](#)
- [41] Sundar Vedula, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. *IEEE TPAMI*, 2005. [1](#), [2](#)
- [42] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a rigid motion prior. In *ICCV*, 2011. [2](#)
- [43] Christoph Vogel, Konrad Schindler, and Stefan Roth. Piecewise rigid scene flow. In *CVPR*, 2013. [1](#), [2](#)
- [44] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a piecewise rigid scene model. *IJCV*, 2015. [2](#)
- [45] Weiyue Wang, Ronald Yu, Qiangui Huang, and Ulrich Neumann. Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In *CVPR*, 2018. [1](#), [2](#)
- [46] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *TOG*, 2019. [2](#)
- [47] Andreas Wedel, Thomas Brox, Tobi Vaudrey, Clemens Rabe, Uwe Franke, and Daniel Cremers. Stereoscopic scene flow computation for 3d motion understanding. *IJCV*, 2011. [2](#)
- [48] Andreas Wedel, Clemens Rabe, Tobi Vaudrey, Thomas Brox, Uwe Franke, and Daniel Cremers. Efficient dense scene flow from sparse or dense stereo data. In *ECCV*, 2008. [2](#)
- [49] Yi Wei, Shaohui Liu, Wang Zhao, and Jiwen Lu. Conditional single-view shape generation for multi-view stereo reconstruction. In *CVPR*, 2019. [2](#)
- [50] Wenxuan Wu, Zhi Yuan Wang, Zhuwen Li, Wei Liu, and Li Fuxin. PointPWC-Net: Cost Volume on Point Clouds for (Self-) Supervised Scene Flow Estimation. In *ECCV*, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [51] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *CVPR*, 2020. [1](#), [2](#)